PROVIGO 410 Business RAID
Users Manual

2^{cnd} Edition (June, 2006)

© Copyright 2004-2006 transtec AG.

transtec AG makes no warranty of any kind in regard to this material, including but not limited to, the implied warranties of merchantability and fitness for particular purpose. transtec AG shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of transtec AG. The information contained in this document is subject to change without notice.

transtec AG shall not be liable for technical or editorial errors or omissions contained herein. The information is provided “as is” without warranty of any kind and is subject to change without notice. The warranties for transtec products are set forth in the express limited warranty statements for such products. Nothing herein should be construed as constituting an additional warranty.

Printed in Germany

PROVIGO 410 Users Manual
2cnd Edition (June 2006)

Table of Contents

Table of Contents **3**
List of Figures **9**
Safety Information **12**
Technical Notes 12
ELECTROSTATIC DISCHARGE 12
RACK WARNINGS 12
System Warnings 13
Symbols on Equipment 13
Introduction **14**
Audience Assumptions 14
About This Guide 14
Getting Familiar with PROVIGO 410 ..... **15**
System Description 15
Key Features 15
Module Locations 19
Status Indicators & System Alerts 21
Status Indicators (LEDs) 21
System Log Notifications 21
Audible Alert 21
Planning your Installation **22**
Host Considerations 22
Host Setup – Cluster 22
Host Setup – Multi-path IO 22
Host Setup – Ethernet Teaming 23
Host Setup – Trade Offs 25
iSCSI Network Considerations 25
Jumbo Frames 25
Teaming 26
PROVIGO 410 Considerations 28
Disks, Disk Groups and Capacity 28
Cache Policy 29
Installing the Hardware **30**
Before you Begin 30
Kit Contents 30
Rack Preparation 31

Rail Preparation and Installation 31
Installing Chassis in Rack 35
Installing Hardware Modules 36
Connecting the Power Cords 37
Configuring the System 39
System Initialization, Network Access and Starting Storage Manager 39
Option 1–Using Pre-configured IP address (recommended) 39
Option 2 - Configuring your own IP address (advanced users) 40
Introducing the PROVIGO 410 GUI 43
Using the Setup Wizard 47
Step 1 – Network Setup 48
Step 2 – Initialize a Storage System 49
Step 3 – Disk Group Setup 51
Step 4 – Virtual Disk Setup 64
Expanding a Snapshot Virtual Disk 75
Step 5 – Host Setup 76
Step 6 – Present Virtual Disks 80
Configuring Options 83
Overview 83
Changing a storage system password 84
Changing the Display Refresh Options 84
Changing System Health Alert Options 85
Setting MPIO Options 86
Configuring the iSNS Client 87
System Upgrade from the GUI 88
Inspection Read 90
Priority 90
Logging 90
Devices to Read 90
Using Inspection Read 90
Options GUI 91
Inspection Read GUI Interface 91
Inspection Read Progress 92
Configuring Notification 93
Using and Maintaining the System 96
System Overview 96
Storage System Properties page 96
Hardware Status 98
System Log 99
ISCSI Statistics 99

System Health Alert 101
Triggers and Corrective Action 102
Temperature Threshold Exceeded 102
Voltage Threshold Exceeded 103
Fan Failure 103
Degraded System Availability 103
Unhealthy RAID Disk Group 104
ECC Memory Error 104
Battery Fault (Cache Battery) 104
Faulty Battery Charger 104
Network Connection Lost 105
Turning the System Health Alert Off 105
System Health Alert Error Notifications: 105
Hot Swap Module Overview 106
Controller Modules 107
Controller Features 107
Controller Features 108
Controller States 109
Manual “Giveback” of IP Address 111
“Fail-back” of Disk Groups 112
Replacing a Failed Controller (Redundant System) 113
Replacing a Failed Controller (Non-redundant System) 115
Upgrading from a Single to Dual Controller 115
Connectors 116
Host Interfaces 117
Teaming Mode 117
Teaming Default 117
Power Supply Modules 117
Power Supply (LEDs) & On/Standby Switch 118
Exchanging a Redundant Power Supply Module 119
Removing All Power from System 122
UPS Vaulting 123
Connection 123
Configuration 124
Messages 124
Notification 124
Cache Battery Modules 125
Hard Disk Modules 126
Rules for Adding Storage 126
Installing Disk Drives or Slot Covers 127
Removing Disk Drives or Slot Covers 128

Hard Drive LEDs 130
Disk Roaming 131
Disk Orphans 131
Appendix A – RAID Concepts & Cache Configurations 133
RAID Levels Defined 133
Cache Settings Defined 134
Appendix B – Using the Command Line Interface 135
Preparing to use the CLI 135
Connecting a Terminal to Access the CLI 135
Accessing the CLI 135
Using the Network Setup Wizard 136
Local/Remote Option 137
Index of CLI Commands 137
Commands 138
System Administration Commands 138
daacl 138
dacomment 139
dadate 139
dahelp 139
dahost 140
dahostname 141
dapasswd 141
dareboot 142
dashutdown 142
dasnmp 143
dastonith 143
dasysalert 144
Network Administration Commands 144
dafcnode 144
daifconfig 144
daisns 145
dans 146
daroute 146
dasetup 146
Storage Administration Commands 147
dadg 147
dagiveback 148
dalun 148
dampio 149
datarget 149

davd 150
System Support and Miscellaneous Commands 151
daaddrbk 151
daconnstat 152
dahwstat 153
dainspread 154
dainventory 155
dakem 156
dasendlog 156
dasetlogmail 157
dasmart 157
dastat 157
daum 158
daversion 159
Appendix C – System Specifications 161
System Enclosure Dimensions 161
Weight 161
Rack Service Clearances 161
Environment 161
Altitude (from sea level) 161
Airflow 162
Electrical Power 162
Input-Voltage Requirements 162
Power Input 162
Power Factor 163
Output Protection 163
Early Power-Off Warning (EPOW) 163
Power Cables 163
Power Control 163
Power Sequencing 163
Auto Restart 163
Acoustic Noise 163
Safety Approvals 163
Electromagnetic Compatibility - Emissions 164
Electromagnetic Compatibility - Immunity 164
Appendix D – Hardware & Software Compatibility List 165
Supported Operating Systems 165
PROVIGO 410 Supported Hard Drives ..... **Error! Bookmark not defined.**
Supported iSCSI Initiators 165
Supported FC HBAs 165

LSI HBAs..... 165
QLogic HBAs..... 165
Appendix E PROVIGO 410 Illustrated Parts List..... 166

List of Figures

Figure 1 PROVIGO 410 Enclosure - Front & Rear View....15
Figure 2 PROVIGO 410 Module Locations - Rear View....19
Figure 3 Configuring Host with Multi-path IO....23
Figure 4 Single host with GE Teaming....24
Figure 5 Teaming without Switch Support....26
Figure 6 Teaming with Switch Support....27
Figure 7 Disk Groups & LUNs....28
Figure 8 PROVIGO 410 - Kit Contents....31
Figure 9 Releasing Chassis Mounting Rail....32
Figure 10 Assembling Rack Brackets to Sliding Rail....33
Figure 11 Installing the Support Rails....34
Figure 12 Installing Mounting Rails to Chassis....35
Figure 13 Installing Chassis in Rack....36
Figure 14 Installing PROVIGO 410 Hard Disk Modules....37
Figure 15 Powering-on the System....38
Figure 16 Login Screen....43
Figure 17 GUI Parts....44
Figure 18 GUI Banner Area....44
Figure 19 Navigation Tabs....45
Figure 20 Setup Wizard....45
Figure 21 Navigation Tree....46
Figure 22 Work Areas....46
Figure 23 System Status Icons....47
Figure 24 Step 1 - Network Setup....48
Figure 25 Step 2 - Initialize a Storage System....49
Figure 26 Step 2 - System Initialization - Date & Time....50
Figure 27 Step 3 - Disk Group Setup....52
Figure 28 Create Disk Group....53
Figure 29 Create a Disk Group entry....54
Figure 30 Disk Group Controller Selections....55
Figure 31 Create Disk Group - Advanced Features....56
Figure 32 Create Disk Group - Progress Page....56
Figure 33 Disk Group Properties....57
Figure 34 RAID Build Progress....59
Figure 35 Build Complete....60
Figure 36 Disk Group Discovery....61
Figure 37 Convert Orphan disk to Spare....62
Figure 38 Disk Group Properties – Expand....62
Figure 39 Expand a Disk Group....63
Figure 40 Expansion Progress....64
Figure 41 Virtual Disk Setup....65

Figure 42 Initial Create VD page 66
Figure 43 Create VD - Selecting Disk Group 66
Figure 44 Create VD - Selecting Cache Policy 67
Figure 45 Create VD - Selecting read-ahead policy 68
Figure 46 Create VD - Selecting VD type 68
Figure 47 Virtual Disk Properties 69
Figure 48 Extending VD Capacity 71
Figure 49 Virtual Disk Snapshot 72
Figure 50 Snapshot Properties 73
Figure 51 Snapshot naming 75
Figure 52 Host Properties Page 76
Figure 53 Adding a Host - Step 1 & 2 77
Figure 54 Adding a Host - Step 3 & 4 78
Figure 55 Host Properties 78
Figure 56 Adding A Host Port 80
Figure 57 Chose a VD to Present 81
Figure 58 VD Presentation 81
Figure 59 Present Virtual Disk Page 82
Figure 60 Present Virtual Disk Advanced Features 82
Figure 61 Options 83
Figure 62 Change Storage System Password 84
Figure 63 Display Refresh Options 85
Figure 64 System Health Alert Options 86
Figure 65 MPIO Options 87
Figure 66 iSNS Configuration 88
Figure 67 Upgrade Manager 89
Figure 68 Options Tab 91
Figure 69 Inspection Read GUI Interface 92
Figure 70 Inspection Read Progress 93
Figure 71 Email Notification 94
Figure 72 Adding Email Addresses 94
Figure 73 Storage System Properties Page 97
Figure 74 Hardware Status PROVIGO 410 98
Figure 75 System Logs 99
Figure 76 iSCSI Statistics Page 100
Figure 77 iSCSI Bandwidth Display 100
Figure 78 System Health Alert Page 102
Figure 79 PROVIGO 410 Front and Rear View 107
Figure 80 PROVIGO 410 Controller Component Locations 109
Figure 81 Active-Takeover warning 110
Figure 82 Active/Takeover Mode 110
Figure 83 IP Address Giveback 111
Figure 84 Disk Group Fail-back 112
Figure 85 Disconnecting Connections from PROVIGO 410 Controller Module 113
Figure 86 Removing the Controller Module 114

Figure 87 PROVIGO 410 System Connector Locations 117
Figure 88 Power Supply Module LEDs and Switch 119
Figure 89 Switching off & Disconnecting the Power Supply Module 120
Figure 90 Unlatching the Power Supply Module 121
Figure 91 Removing the Power Supply Module 121
Figure 92 Installing a PROVIGO 410 Hard Disk Module or Slot Cover 127
Figure 93 Opening the Handle of a PROVIGO 410 Disk Drive Module 129
Figure 94 Removing a PROVIGO 410 Disk Drive Module 129
Figure 95 PROVIGO 410 Disk Drive Module Light Location 130

Safety Information

Technical Notes

- **WARNING:** Only authorized technicians should attempt to repair this equipment. All trouble shooting and repair procedures allow only module replacement. Because of the complexity of the individual boards and sub-assembles, no one should attempt to make repairs at the component level or make modifications to any printed wiring board. Improper repairs can create a safety hazard.
- **WARNING:** To reduce the risk of personal injury from electric shock and hazardous energy levels, do not exceed the level of repairs specified in these procedures.
- **CAUTION:** The system is designed to be electrically grounded. To ensure proper operation, plug the AC power cord into a properly grounded AC outlet only.

ELECTROSTATIC DISCHARGE

Electrostatic discharge (ESD) can damage static sensitive devices or micro circuitry. Proper packaging and grounding techniques are required to prevent damage.

- Transport products in static safe containers such as conductive tubes, bags or boxes.
- Keep electrostatic-sensitive parts in their containers until they arrive at a static free work area.
- Use a wrist strap connected to the work surface as well as properly grounded tools and equipment
- Keep the area free of nonconductive materials such as ordinary plastic assembly aids and foam packing.
- Ensure proper grounding before touching a static-sensitive component or assembly.
- Avoid touching pins, leads, or circuitry.
- Always place drives with printed circuit board (PCB) assembly-side down.

RACK WARNINGS

If you plan to rack mount the PROVIGO 410, please follow the rack manufacturer's safety instructions. Install the enclosure only in a rack that has been properly secured in an area with suitable environmental conditions.

- **CAUTION:** Have someone assist you during physical installation.
- **CAUTION:** To properly ventilate the system, you must provide at least 7.6 cm of clearance at the front and back of the server.

- **WARNING:** To reduce the risk of personal injury or damage to equipment, always ensure that the rack is adequately stabilized prior to extending a component outside the rack. A rack may become unstable if more than one component is extended. Extend only one at a time.

System Warnings

- **WARNING:** To reduce the risk of personal injury from hot surfaces, allow the hot-plug disk modules and other system modules to cool before touching them.
- **WARNING:** To reduce the risk of electric shock or damage to the equipment:
  - Do not disable the power cord grounding plug. The grounding plug is an important safety feature.
  - Ensure the power cord is inserted into a grounded electrical outlet that is easily accessible at all times.
  - Unplug the power cord from the power supply module to shut off power to the equipment
- **CAUTION:** Protect the storage system from power fluctuations and temporary power interruptions with a regulating uninterruptible power supply (UPS). This device protects the hardware from damage caused by power surges and voltage spikes and keeps the system operational during a power failure.
- **CAUTION:** The storage system must always be operated with all hot plug modules installed or slot covers in place to ensure proper cooling.

Symbols on Equipment

WARNING: Hot Surface or Component - To reduce risk of injury from a hot component; allow the surface to cool before touching.

WARNING: Electric Shock Hazard – To reduce risk of injury from electric shock hazards, do not open this component.

WARNING: Contains No User or Field Serviceable Parts – To reduce the risk of injury from electric shock hazards, do not open this component.

WARNING: Host Interface Only- To reduce the risk of electric shock, fire or damage to equipment, do not plug telephone or telecommunications connectors into this receptacle

Introduction

Audience Assumptions

This guide assumes that you are a service technician or network administrator familiar with computer hardware, data storage and network administration terminology and tasks. It also assumes that you have a basic knowledge of networking and RAID technology concepts.

About This Guide

This guide provides step by step instructions for installation, use and maintenance of the PROVIGO 410 Intelligent Storage Arrays. This guide is generally organized as follows:

Getting Familiar with the System – System description and features.
Planning Your Installation – What you need to do before you start.
Installing the Hardware - Instructions on installing all hardware modules and connections.
Storage Configuration – Initializing your system, setting up your system on the network, configuring your storage, and connecting to hosts.
Using & Maintaining Your System – Ongoing usage and maintenance information provided in a reference format and organized by major system component.

Getting Familiar with PROVIGO 410

System Description

The PROVIGO 410 is a low cost, high availability RAID array. Unlike most arrays designed to support a single controller with proprietary architectures, this storage platform can be configured as a single or dual controller system. They are built around a flexible and open architecture utilizing dual; active-active Pentium based storage controllers, and a custom tuned Linux operating system. PROVIGO 410 is offered in a 3U rack mount form factor with the ability to scale to 11.25 TB (or higher as drive capacity grows beyond 750GB) of raw storage capacity. The system offers DAS and SAN functionality and provides both iSCSI and optional FC host connectivity.

The PROVIGO 410 supports SATA-I and SATA-II disk drives and with the SAS connector, the PROVIGO 410 provides an additional SAS connection on each controller for use with a SAS-to-SATA expansion chassis.

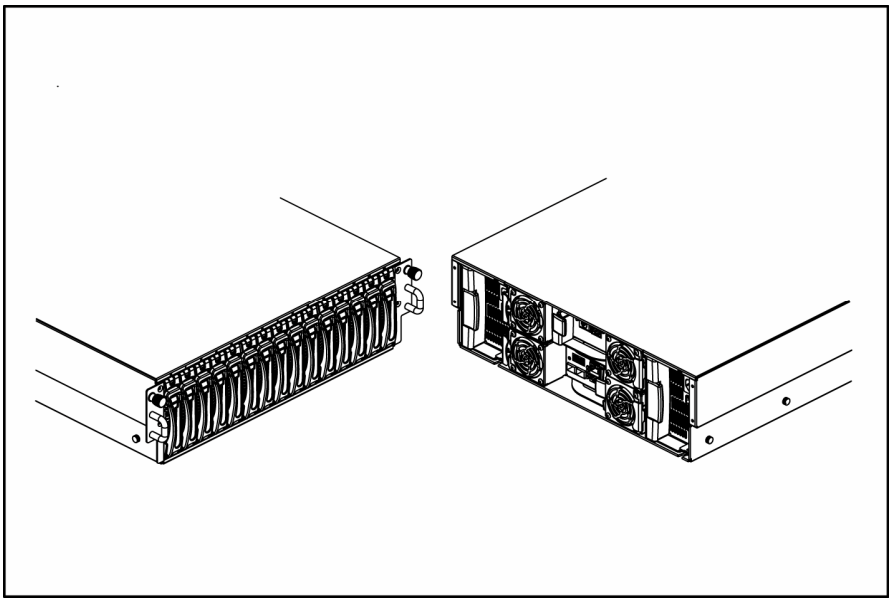

Figure 1 PROVIGO 410 Enclosure - Front & Rear View

Key Features

Zero Down-Time Architecture - In the PROVIGO 410 every active component, including drives, controllers, fan trays, and power supplies, are hot swappable without interruption in service or loss of data. Its active-active dual controllers, mirrored and

battery backed cache and path fail over capabilities give the PROVIGO 410 system fault tolerance.

Robust Enclosure - Both the enclosure and the drive carriers are designed for operation without IO errors in various temperature and vibration conditions. The whole enclosure is built in one piece. Extensive thermal and vibration testing has been conducted on the enclosure itself and on the rack with stacked chassis.

Passive Mid-plane - PROVIGO 410 is built around a completely passive mid-plane. There are no components other than connectors in the mid-plane, a feature not commonly found in other storage systems of its kind and price point.

In the PROVIGO 410, special attention has been taken to recover the data path failure due to signals stuck at zero or one. In addition, each active component such as the drives and controllers monitor the incoming current and shut down power in case of a short-circuit.

Active-Active Controllers & Transparent Fail-over - Fast throughput and data availability are paramount for dedicated storage systems. This is why PROVIGO 410 is designed with active-active controllers to improve both performance and data availability. Each disk group has an assigned "primary" controller to take advantage of the fast route. So, optimal performance is achieved when traffic goes through both controllers to their associated disk groups. For availability, both controllers can serve any disk groups within the system. When there is a controller failure, the IP address and disk groups assigned to the failed controller will automatically "migrate" to the surviving controller without user intervention.

"Hardware" In-box Clustering - In PROVIGO 410, the relationship between two controllers is tight. Each controller has the precise status of the other, which results in quick and accurate actions when failure happens.

Multi-Function Flexibility – The PROVIGO 410 is offered in a 3U rack mount form factor with the ability to scale to 11.25 TB (or higher as drive capacity grows beyond 750GB) of raw storage capacity. The system offers DAS and SAN functionality and provides iSCSI or FC host connectivity within the same box. This addresses diverse segments in the storage space without the need for multiple specialized platforms.

Mirrored Cache - With dual controllers, the "auto cache" option puts a mirrored copy of the write cache on the redundant controller. The data is fully protected. If there is a failure in any controller, the system flushes the mirrored copy and automatically switches to the "write-through" setting. This is the best compromise between performance and protection.

Multi-Path IO and Cache Coherency - In PROVIGO 410, the system makes no assumption regarding the behavior of the hosts since they can switch from one path to

the other. This is normally referred as “multi-path IO” – a technology to ensure the end-to-end high availability for path failure. The patent-pending technology supports path failure with our cache coherency features. In the extreme case, the host can send commands to either controller randomly. Uninterrupted I/O access in the event of a controller failure requires 3rd part software (e.g. Symantec Storage Foundation Suite) for automated FC host path failover within Microsoft Windows host servers.

Dedicated SATA/SAS Channels - PROVIGO 410’s multi-path design provides a dedicated SATA/SAS drive channel from each controller. The elimination of inter-drive dependencies enhances performance. The independent controller access provides additional data availability. The patent-pending dual-access scheme provides fair and robust access from each controller.

Disk Roaming - In PROVIGO 410, the concept of disk groups is introduced, which represents a set of drives that share similar characteristics. The disk group configuration, including RAID, and volume, roams with the drives. Once inserted, the disk group can be discovered and the service can start with a mouse click. Within the enclosure, the drives can also be moved around. The roaming operation can be carried out while the system is on-line. Users can take advantage of this flexibility to arrange their storage for different on-line and off-line tasks. This feature provides a basic and inexpensive way to upgrade a system and retain your data or move drives to another system for basic system recovery.

Capacity Expansion - The capacity expansion feature simplifies the task of adding additional storage capacity. The user can expand the LUN capacity on-line within the disk group, or expand the disk group by adding disks from the pool of global spare disks. Adding disks to a disk group may require multiple disks.

Software RAID - The PROVIGO 410 architecture focuses on the critical performance elements of protocol processing and data movement. The powerful Intel P4 processor and server chipset (high-speed front-side bus and DDR) present a data mover architecture that is superior over embedded processor solutions found in traditional storage arrays.

Traditional embedded solutions tend to be a generation behind in memory technology and require special chips to enhance performance. With a dedicated Pentium in our array, there is little benefit of using expensive “hardware” RAID. “Hardware” RAID is important in general-purpose servers to off-load the CPU. In an external array with dedicated controllers and memory, “hardware” RAID only adds cost.

Global Spares - In PROVIGO 410, any drives not assigned to a disk group are automatically available as global hot spares. If a drive fails, the RAID controller can automatically rebuild its data on the spare drive without requiring intervention by the administrator. The rebuild operation occurs in the background while the controller processes normal operations.

LUN Masking and CHAP Authentication - Enterprise data must be secured and closely guarded. In PROVIGO 410, LUNs are only exposed to the hosts that have authority to access them. This is generally referred to as the LUN masking. In addition, CHAP authentication is used for each iSCSI session login. All management accesses are also protected with the access control list, which filters out the machines that are not on the list.

Point in Time Snapshots – PROVIGO 410 supports manually triggered, point in time (PIT), full copy snapshots on virtual disks. Snapshots enable you to present the read-only snapshot to a second server without interrupting operation on the original VD. This is useful when you want to back up your data without effecting operation, or when you want to provide some clients with read-only access to data on a specific Virtual Disk.

Simple, Intuitive Configuration - The web based configuration utility RAIDconn has been designed to be intuitive and extremely easy to use and requires little to no training. With RAIDconn, a system administrator can set up a system in 6 simple steps.

The utility allows you to:

- Assign unique IP address to each controller
- Select and name multiple drive groups and establish different RAID levels for each group.
- Create and size virtual drives for each drive group. The system supports up to 256 LUNs.
- Create a list of available hosts with WWN's and IP addresses
- Create point-in-time snapshots
- Establish access rights to each virtual drive with the option to enable powerful security policies to each virtual disk. Add capacity, check system status, set email alerts in the event of system problems.

System Management and Full Monitoring Capability - transtec is committed to developing storage arrays that reduce administrative complexity and management costs. PROVIGO 410 comes with powerful remote management software that dramatically simplifies setup and monitoring of storage volumes. The robust self-monitoring technology and automatic notification service maximizes storage availability while minimizing the pressure on system administrators.

In the event of a hardware failure, PROVIGO 410 can automatically notify you via email. The user interface also provides monitoring for each hardware module. The log information is kept in the persistent memory, which can be recovered after reboot. For remote diagnostics, the administrator can dump the whole log to a remote email address with a single mouse click.

On-Line Firmware Update - Using PROVIGO 410, there is no need to take your system down when upgrading the firmware. As part of the high-availability design, PROVIGO 410 keeps mirrored copies of the RAID firmware KDI in flash on each controller. You can upgrade the kernel on-line with no disruption of service.

Module Locations

The PROVIGO 410 module locations and their associated slot numbers are identified in the diagram below (See **Figure 2**)**.** When more than one module of a type is installed, for example, two controller modules, each is identified by a slot number.

The slot number assigned to each module corresponds to the position in which it is installed in the enclosure.

① System Enclosure
② Power Supply Module (PSM) – Slot #0
③ Power Supply Module (PSM) – Slot #1
④ Controller Module (CM) – Slot #1
⑤ Controller Module (CM) – Slot #0
- Battery module (in the CM)
- Fans (in the CM)
- Hard Disk Modules (HDM)

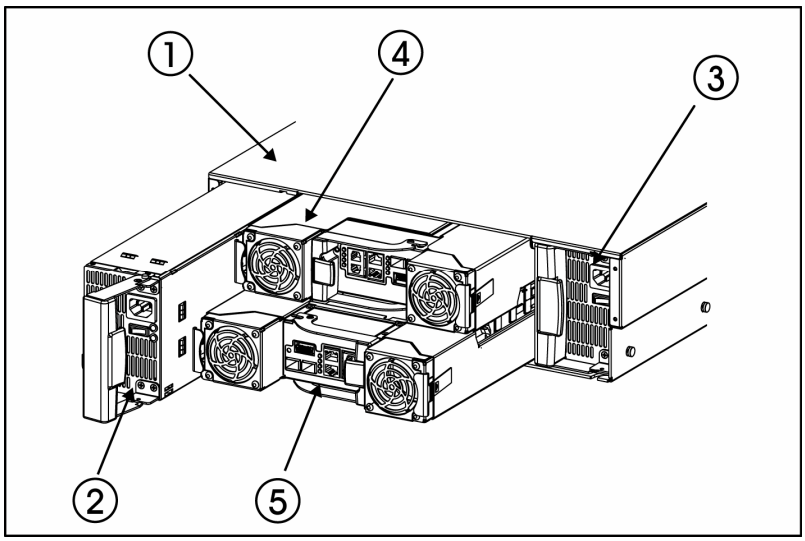

Figure 2 PROVIGO 410 Module Locations - Rear View

Note: This view shows modules that can be inserted and removed from the back of the enclosure. The following applies when configuring a system:

- An enclosure can have one or two controller modules. When a controller slot does not contain a controller module, a controller slot cover must be installed to maintain proper airflow.
- Every enclosure has either one or two power supply modules. When a power supply slot does not contain a power supply module, a power supply slot cover needs to be installed to maintain proper airflow.
- Disk drive modules are numbered 0 through 14 from left to right (See **Figure 1**)

Status Indicators & System Alerts

The PROVIGO 410 is equipped with a full complement of status indicators and system alerts.

Status Indicators (LEDs)

PROVIGO 410 has multi-color LEDs that located on all the removable modules. These provide system administrators visual cues to the status of all major components.

The LEDs are located on:

- Controllers
- Fibre Channel and Ethernet Interface Ports
- Power Supplies
- Hard Drives

System Log Notifications

PROVIGO 410 also captures status and health alerts in the system log. These alerts are viewable by the administrator in the PROVIGO 410 RAIDconn. The alerts can be viewed through the “System Log” or “System Health Alert” sub-menus. For proactive notification, the administrator can set up to 8 email addresses to receive system log notifications.

Audible Alert

The PROVIGO 410 has an audible alert that is part of the system startup and initialization sequence. When the system is started for the first time or in cases when the system is improperly shutdown, the system waits 2 minutes to initialize and run a full system check. The system makes an audible “ticking” sound as it counts down through the initialization sequence.

Note: If the system is restarted during the initialization sequence, the process starts over from the beginning. Please allow the system to finish its initialization sequence to avoid system startup delays.

Planning your Installation

When deploying the PROVIGO 410 in a compute environment that consists of hosts, network devices and storage careful planning will provide the best solution for your deployment. This section will review elements that impact peak performance, availability, cost and ease of configuration.

Appendix A provides examples of deployments that optimize specific characteristics toward application requirements and environments.

Host Considerations

Host considerations include the system limitations, connection types, data path configuration and other characteristics. In this section we will look at a few of the best practices in use today to optimize the host.

Host Setup – Cluster

By clustering multiple nodes, there is no single point of failure at the host. Here, a "host" consists of two or more nodes. Each node has its own network path to PROVIGO 410. When a failure happens in the active node, fail-over happens to transition the service to the surviving node. The surviving node will re-establish the access to the storage. Microsoft Cluster Services (MSCS) and Linux cluster software are two examples of clustering software.

Host Setup – Multi-path IO

In certain environments, a single host is used. This is normally for larger machines (e.g., multi-core SMP, UNIX servers), which have built-in robustness and the cost of clustering is un-prohibited.

The host must have two GE ports to provide redundant path to PROVIGO 410. When a failure happens in one of the paths, the host can switch to another. The host must be equipped with the multi-path IO software. As an example, Microsoft iSCSI initiator versions 2.0 and higher have MPOI support built into the initiator.

The MPIO does the path fail-over and load balancing at SCSI command level. For example, you can have the read commands always come from one path and write commands from another. Or, you can alternate the commands to each route.

MPIO is a storage protocol. It has the notion of dual-path and shared storage. When the host goes through two separated paths to reach the storage, the MPIO can identify the shared LUNs by its unique IDs. In other words, the same LUN will not show up twice in the host via separate paths.

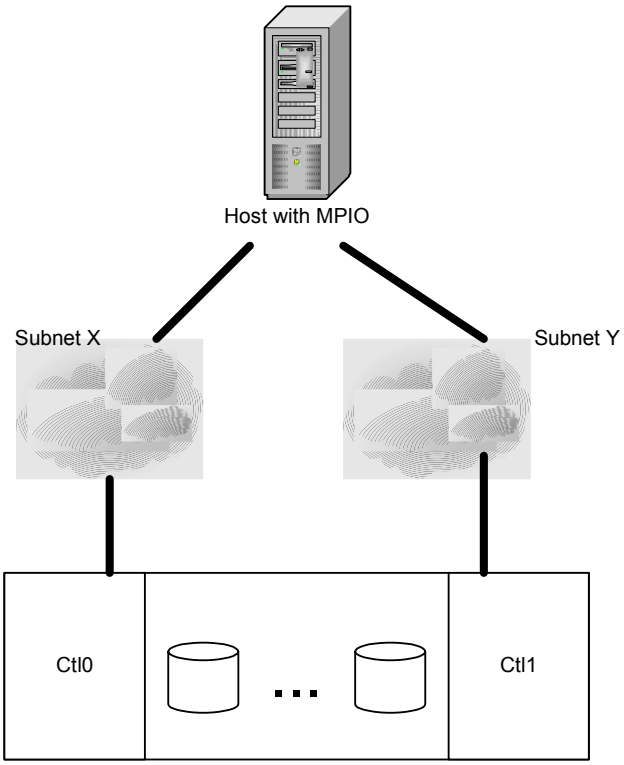

Figure 3 Configuring Host with Multi-path IO

Host Setup – Ethernet Teaming

Ethernet Teaming is an alternative to MPIO. The host still utilizes two GE ports, but in this case instead of relying on the MPIO software, both GE ports are teamed together and the driver does the work.

GE teaming provides the capability of network path fail-over. When one path is down, the traffic can be routed to the redundant path.

GE teaming also does limited load balancing. The load balancing granularity is per TCP session or (source IP, destination IP) set. Specifically, the teaming algorithm directs the same TCP session data through the same path. Different TCP sessions with different IP destination and source sets use different paths. For traffic between the same host and the same target, it is impossible to use two paths. So, compared with 2Gbps FC port,

dual GE ports do not achieve equivalent speed as FC due to the teaming algorithm mentioned above.

Since teaming is a network protocol, it has no notion of shared storage. The fail-over and load balancing only happen between network end points. In other words, there is no shared storage concept. To deal with this limitation, during fail-over, the surviving controller must do IP address take-over to re-create the same end point. This is not required for clustering and MPIO.

For the same reason, load balancing in teaming terminates at the network end point. It cannot take advantage of the controller with different IP address, i.e., cannot take advantage of the active-active controller mode (See **Figure 4**).

Because teaming is a network protocol, everything is taken care of inside the network. There is no additional configuration required and everything is automatic. For the same reason, there is no way to fine-tuned fail-over or load-balancing as you can with MPIO.

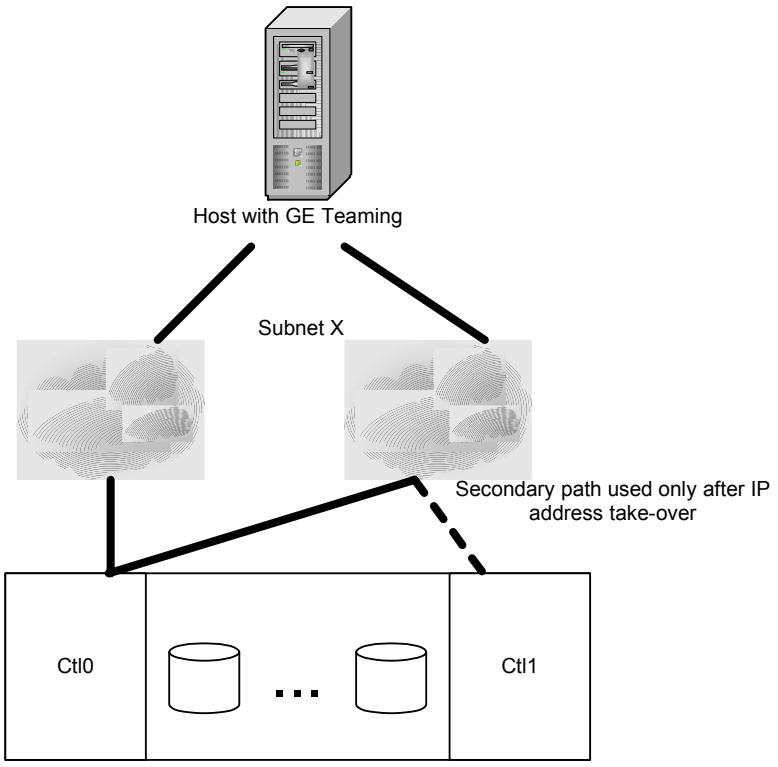

Figure 4 Single host with GE Teaming

Host Setup – Trade Offs

PROVIGO 410 supports all the above host configurations. Here is a summary to help plan based on High-Availability (HA), cost of deployment, performance, and ease-of-configuration.

HA. The best HA is achieved with host clustering. This provides a fully protected environment from the host, network, and storage. There is indeed no single point of failure.

Cost. Clustering is expensive. Redundant host hardware is needed, and you also have the cost of the clustering software. Each cluster node typically requires a separate software license. The MPIO saves some cost by not incurring the software expense. GE teaming is the least expensive because teaming comes free with chip and driver.

Performance. Clustering operations have minimal impact on the host resource. MPIO, on the other hand, requires host CPU to direct SCSI commands thus impacting host performance. GE teaming is also done in software. Teaming requires resources to check each packet's destination and source IP address. GE teaming impacts both host and PROVIGO 410. Therefore, teaming has the biggest impact to the overall processor performance but may be appropriate for large data transfers where the overhead is offset by the increased data throughput.

Ease-of-Configuration. Clustering can be difficult to set up. MPIO requires special software and associated tuning. Teaming is effortless. From the network perspective, both clustering and MPIO require separated IP subnets. This needs additional network planning. Teaming, again, is effortless because multiple GE ports are teamed together with a single IP address. So, only a single subnet is needed. The PROVIGO 410 ships with Teaming enabled by default.

iSCSI Network Considerations

In this section we will look at two of the factors that effect performance on an iSCSI storage network. This section does not attempt to address the performance characteristics of switches or complex network topologies, but rather focus on details that specifically impact data throughput.

Jumbo Frames

Both the hosts and PROVIGO 410 can support Ethernet Jumbo frames. Jumbo frames allow larger packet sizes from the standard 1500 bytes to 9000 bytes. This improves network efficiency because more data is transported with the same network overhead. Jumbo frames further improve the processing efficiency at both the host and PROVIGO 410, because less processing is needed for the same set of data.

Note that not all GE switches support jumbo frames, especially in low-end. PROVIGO 410 is shipped without jumbo frames enabled by default to accommodate this. To improve the performance, turn on jumbo frames if your network can support it.

Teaming

Each PROVIGO 410 controller has two Broadcom GE NIC adapters. During boot up, Broadcom Advance Server Program (BASP) is started to enable teaming.

BASP is a Broadcom intermediate software driver for Linux that provides load-balancing, fault-tolerance, and VLAN features. These features are provided by creating teams (virtual adapters) that consist of multiple NIC interfaces. Teams in the PROVIGO 410 controllers consist of the two NIC interfaces. Both interfaces in a team participate in Load-balancing operations by sending and receiving a portion of the total traffic. Either interface will take over in the event that one of the interfaces lose their links. Load-balancing and fault-tolerance features will work with the Broadcom NIC adapters.

In the PROVIGO 410 controller, the VLAN ID is designated as da0. This is the address that is configured during the network configuration step in the system set up procedure. The assigned address for each controller is saved in /etc/niconf-0 for controller 0 and /etc/niconf-1 for controller 1. Ports eth0 and eth1 on each controller are up running in slave multicast mode.

Not all GE switches support teaming, especially in low-end. Teaming is different from jumbo frames though in that teaming can function without the switch support in certain configurations. Following are two example configurations where the end points perform the teaming and switch support for teaming is not required.

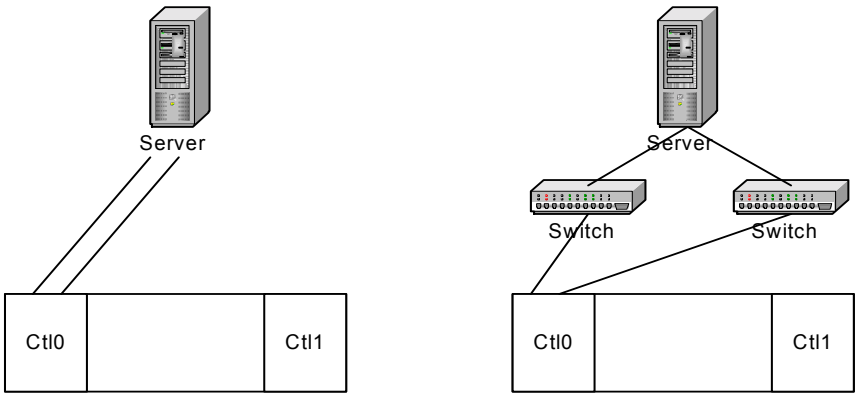

Figure 5 Teaming without Switch Support

In the following configuration, switch teaming is needed because one side of the teaming end point is on the switch.

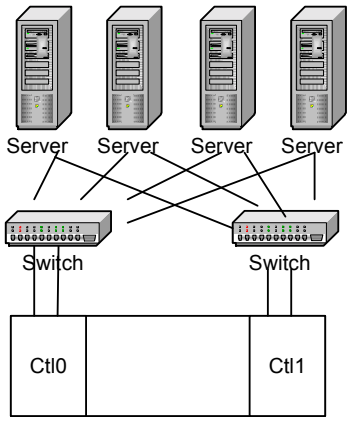

Figure 6 Teaming with Switch Support

PROVIGO 410 Considerations

In addition to optimizing the hosts and network for your needs, there are several areas to consider specific to the configuration and use of the PROVIGO 410. In this section we will discuss the impact of basic choices in disks to the practical implications of the caching policy.

Disks, Disk Groups and Capacity

A disk is a mechanical device. Delivering data is slower if the head has to move to reach the data. Head movement takes multiple milliseconds. For this reason you will see lower throughput with random access tests. Random access increases the disk head movement. You will see the best throughput with sequential access, where head movement is minimized.

The 'disk group' concept in PROVIGO 410 combines (or groups) a number of disks together in a particular RAID set. You may have up to four disk groups in a single chassis. Virtual disks, or LUNs, can be created inside the disk group. If multiple LUNs inside the disk group are accessed at the same time, you will see a reduction in performance due to the head movement required to move from one LUN to others. This even happens with sequential access.

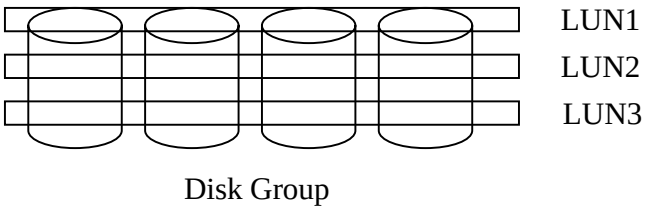

Figure 7 Disk Groups & LUNs

As a result, the best performance numbers are achieved by giving each host its own disk group to reduce head movement.

Giving each host its own disk group means more disk groups are needed. More disk groups require more space for parity since each disk group has its own protection. For example; if RAID-5 is used for four disk groups, an equivalent of four drives are consumed for parity.

Cache Policy

Use a write-back (WB) caching policy for peak performance. Write-back caching stores data in the controller cache before writing to the disk. This approach reduces the number of disk write operations. However, you should realize that Write-back caching is vulnerable to controller failure, see “**Cache Settings Defined**” on page **134** for more information.

Write-through (WT) caching provides for data integrity by writing to the disk and cache at the same time. Write-through caching does not improve performance of write operations, but makes subsequent read operations faster when the requested data is available in cache.

The AUTO cache policy is the best comprise between performance and protection. It automatically switches between mirrored cache with a Write-back policy when two controllers are present and Write-through caching when only one controller is available. The AUTO cache does carry the overhead for mirroring data to the other controller, so the performance is less than Write-back on a single controller.

Installing the Hardware

Before you Begin

Before you setup your PROVIGO 410, consider the following:

Thoroughly research and establish an installation and configuration plan for your specific network environment. You should also plan how you want to configure your storage. (See "**Planning your Installation**" starting on page **22** for an introduction to configuration planning.)

The system comes pre-set with two default IP addresses. The simple online configuration procedures will allow you to set-up your storage on the network using these IP addresses. If your network administrator requires private IP addresses, they should provide the following information so that you can change the default configuration:

- IP Address
- Gateway address
- Net mask address
- DNS server information for each controller.

Kit Contents

Make sure you have all the components that shipped with your system. The PROVIGO 410 System ships with the following:

- **Front Bezel** – the lockable front bezel helps limit access to hard drives and minimizes risk of unintentional removal
- **Chassis** – The chassis ships with 1 or 2 Controller Modules (CM) installed. Each slot not populated with a CM has a slot cover for proper cooling. The chassis also ships with 1 or 2 Power Supply Modules (PSM). Each slot not populated with a PSM has a slot cover for proper cooling.
- **Hard Disks** – 15 Hard Disk Modules (HDM) are shipped with your system. Each HDM consists of a hot plug canister, a dual-port connector and, depending upon chosen configuration, a SATA hard drive
- **Power Cords** – The system ships with 1 or 2 power cords.
- **Serial Cable** – The system ships with a serial (RS232) null modem cable for use with the management port on the system.
- **Rail kit** – For installing the chassis in a rack.

- **Documentation CD**.

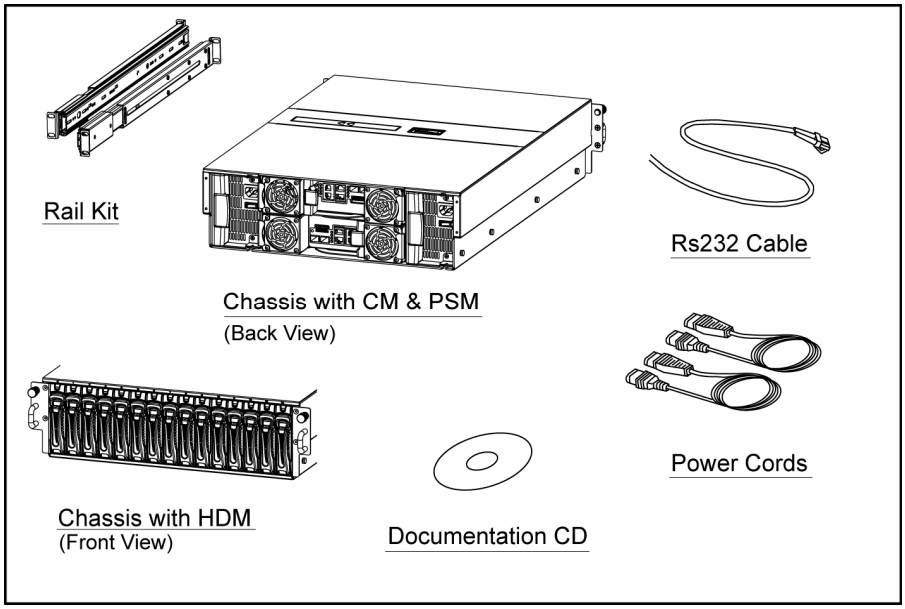

Figure 8 PROVIGO 410 - Kit Contents

Rack Preparation

1. Review the rack manufacturer's documentation to become familiar with their installation requirements.
2. Locate the two support rails that are supplied with the PROVIGO 410. Have them ready for use in the rack installation.
3. If you are installing more than one PROVIGO 410 into the rack, start at the lowest available position, and work upward.
4. You may need to remove the rack's power distribution unit (PDU) before you install the support rails. Refer to rack manufacturer's instructions for rack preparation.

Rail Preparation and Installation

1. Locate the rail kit that ships with PROVIGO 410. The rail kit ships with 3 pairs of parts, 2 small front brackets, 2 longer rear brackets and 2 sliding rails. It also comes with a small packet of assembly hardware.
2. Start with one of the sliding rails. Lay it flat with the smooth side up. Pull out the bottom rail a little bit. It will slide out from only one side of the rail. A small white button is visible as you slide it out.

3. Press the white button away from the rail with your thumb and pull this piece out completely from the rail.
4. Put this part aside. It is the chassis mounting rail which will be snapped onto the chassis in a later step (See **Figure 9**).

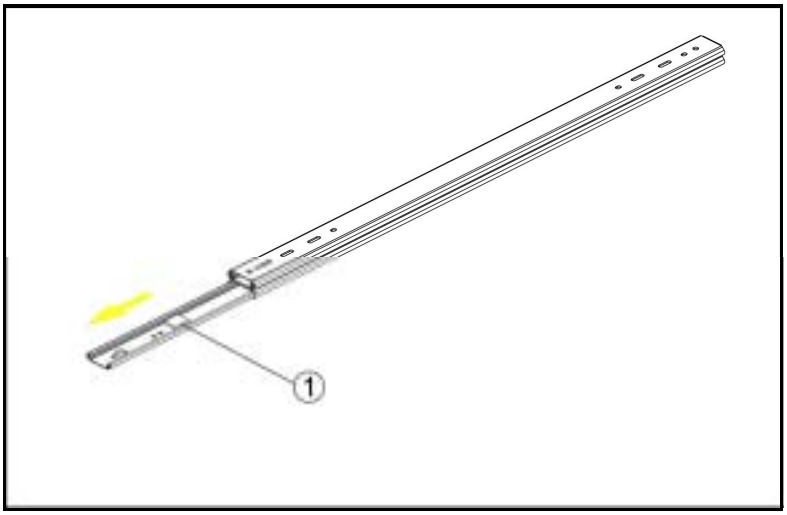

Figure 9 Releasing Chassis Mounting Rail

5. Flip over the sliding rail and extend until partly open. This side of the rail will show the ball bearing mechanism. Take one of the rear brackets and slide the bracket onto the slide rail opposite from where you removed the last piece.
6. Use the screws to assemble the rear bracket to the sliding rail. There are many holes on the bracket and rail. You can adjust the position of the bracket using these holes. Do not tighten the screws fully (See **Figure 10**).

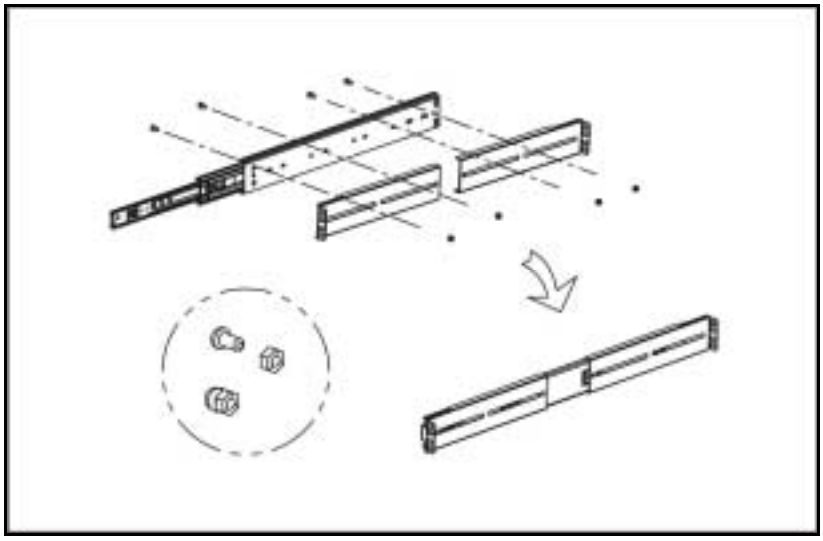

Figure 10 Assembling Rack Brackets to Sliding Rail

7. Take one of the front brackets (smaller bracket) and slide onto the other end of the rail. Assemble the front bracket like you did the rear. Make sure the bracket is on the outer part of the slide rail, not the inner sliding mechanism. It should be assembled to the same side as the rear bracket. You will need to slide the rail until you are able to align and access the holes on the rail. Use the screws to assemble the front bracket to the sliding rail. Use the large hole on the underside of the rail to access and tighten the screws (See **Figure 11***)*.

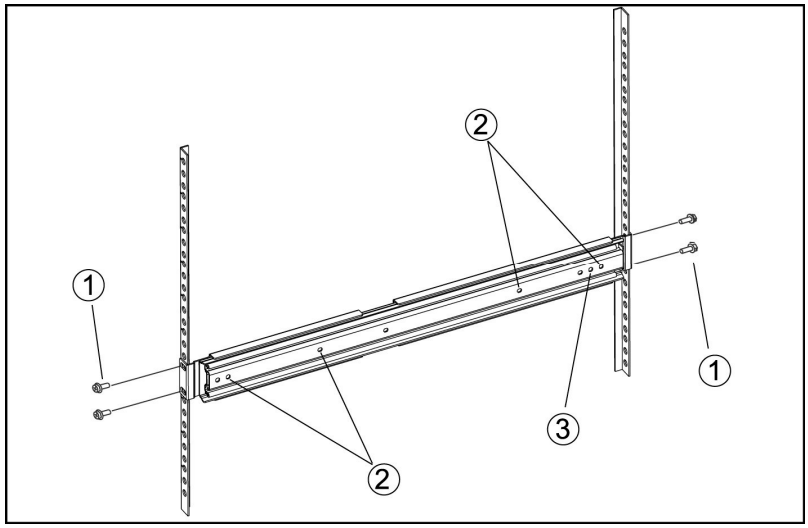

Figure 11 Installing the Support Rails

8. Each PROVIGO 410 enclosure is 3 EIA units high. Locate the EIA markings on the rack and decide where you are going to install the support rails (See **Figure 11**).
9. Hold the sliding rail and bracket assembly you just worked on so that it is horizontal, and position it so that the rack screws ① align on the rack. Adjust the length by extending the rear bracket. Slide the nut clips onto each hole of the rack and tighten the rack screws① to the nut clips.
10. Tighten the rail screws ② so that the rail is held firmly in place on the rack.
11. Verify that the support rails are horizontal. If not, remove them from the rack and repeat rail installation steps 7-8.
12. Repeat steps 2 through 11 for the second rack rail.

Note:

- The ends of the rail should be positioned outside the rack supports.
- The holes noted ③ are not used to mount the support rails. They are used to attach the enclosure to the support rails in a later step.

Installing Chassis in Rack

Warning:

- A stabilizer must be correctly attached to the bottom, front of the rack to prevent the rack from tipping forward while the enclosure is being removed or installed. Do not remove or install any enclosure if a stabilizer is not attached to the rack.
- Do not attempt to lift the enclosure by yourself. Ask another person for aid.
- Do not use the handles from the controller or power supply modules to carry the enclosure. These handles are not intended to support the weight of the enclosure.

1. Take the chassis mounting rail removed from the sliding rail in the previous section in Step 3. Attach the mounting rail to the chassis. Make sure that the end with the white button is to the front of PROVIGO 410 (See **Figure 12**).

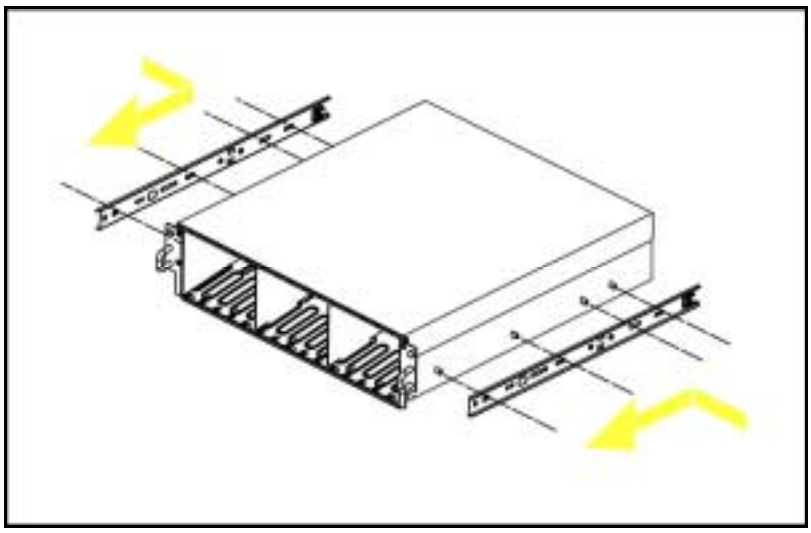

Figure 12 Installing Mounting Rails to Chassis

2. Before installing the chassis in the rack, remove the controller modules, power supply modules and ensure that there are no hard disk modules installed in the chassis. This will make the chassis easier to lift and maneuver for installation in the rack (See **Figure 12**).

3. Stand at the front of the rack and with help from another person, place the mounting rail at the back of the enclosure onto the sliding support rails, and then slide the enclosure into the rack.
4. Install the two front mounting screws ① (See **Figure 13**) into the top nut clips you placed on the rack and tighten them.

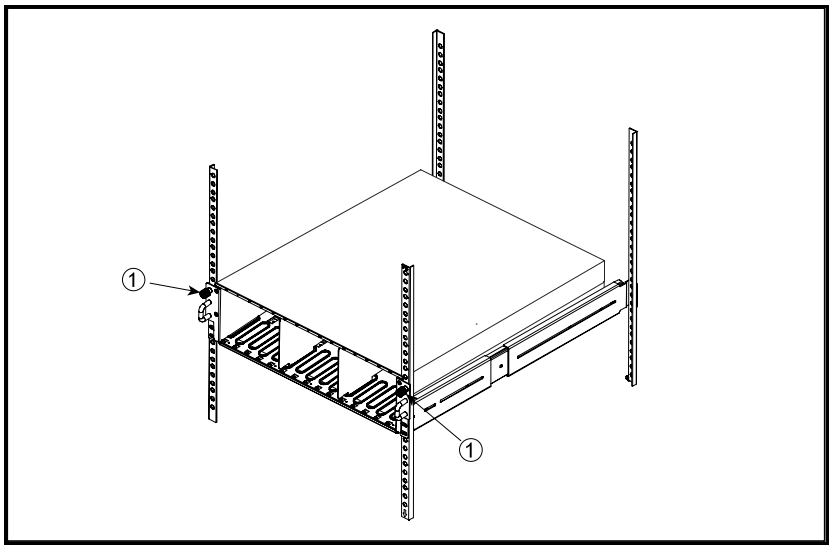

Figure 13 Installing Chassis in Rack

5. At the back of the chassis (See **Figure 11**) tighten the two screws on each side of the chassis③.
6. Proceed with “Installing the Hardware Modules” in the next section.

Installing Hardware Modules

1. Unpack the hard disk modules. Ensure that the handle ① is fully open on the disk drive module that you are going to install. With one hand supporting the base of the module and the other hand holding the handle, insert the module and push it into the slot. When the handle touches the front surface of the subsystem, the module will stop. (The module is not yet fully inserted.) While continuing to push the disk drive module into the slot, slowly raise and close the handle ① until it snaps into place. Follow the same procedure for all remaining drives.

Note: There must be a minimum of 3 drives to configure a disk group in the system. (See "**Rules for Adding Storage**" on page **126**)

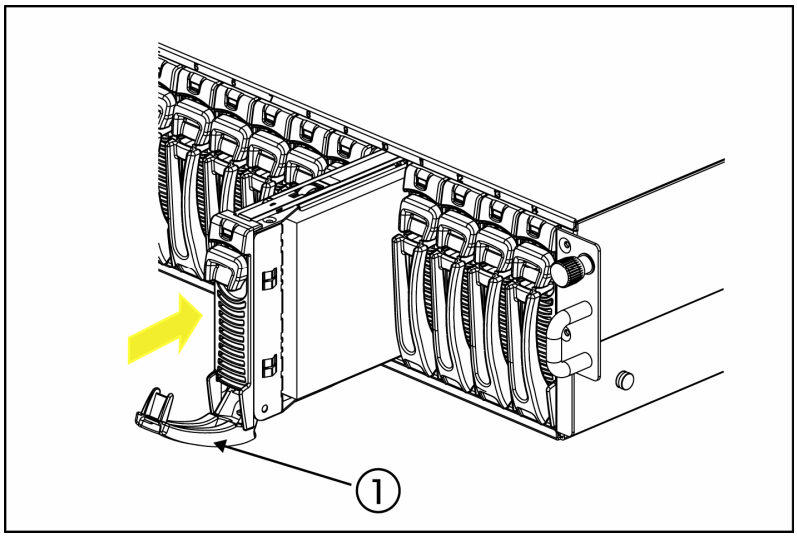

Figure 14 Installing PROVIGO 410 Hard Disk Modules

1. Reinstall the controller modules and power supply modules that shipped with the system. Ensure they are fully inserted and snapped into place.
2. Use Gigabit Ethernet cables to connect the PROVIGO 410 to your network.

Connecting the Power Cords

CAUTION:

- PROVIGO 410 is equipped with a power cable and plug for user safety. Use only the cable provided with the system. Ensure that the electrical outlet is correctly grounded to avoid electric shock.
- Do not try to configure and install this product during an electrical storm.

1. At the back of the enclosure, plug the power cords ② to the power supplies. Plug the other end of the power cord into the power source of the rack (See **Figure 15**). Make sure the power source is fully grounded. It is recommended that you always use a regulating uninterruptible power supply (UPS) to protect your system.

Note: Do not turn on standby switch on either power supply until you are asked to do so a later section. This will facilitate the initial network configuration.

2. The physical hardware installation is now complete. Go to the next section to prepare your host systems.

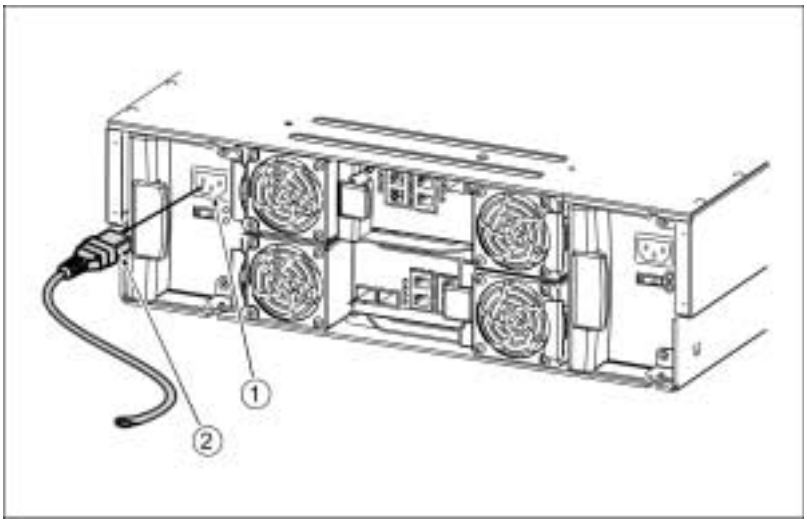

Figure 15 Powering-on the System

Configuring the System

System Initialization, Network Access and Starting Storage Manager

Once the physical hardware components are installed, you are ready to initialize the system and start configuring your storage. First you will need to connect and set up PROVIGO 410 on your network. There are two options for setting up the PROVIGO 410 on your network.

Option 1–Using Pre-configured IP address (recommended)

Each PROVIGO 410 comes pre-configured with an IP addresses for each controller. Once you have connected your PROVIGO 410 to your network and initiated the system, you can use the web based PROVIGO 410 RAIDconn included with your system to configure your storage.

Using the PROVIGO 410 RAIDconn and default IP addresses instead of the CLI is recommended, especially if this is your first time setting up a PROVIGO 410. The PROVIGO 410 RAIDconn is designed to get your system up and running in six simple steps.

1. Power-on PROVIGO 410, using the DC power/standby switch-on each power supply. The green DC power light on the power supply will turn on and stay on. The green activity light on each installed disk drive should be on. When the system is powered on, the embedded RAID kernel boots automatically.

Note: When PROVIGO 410 is started for the first time, the system may take 1 to 2 minutes to initialize as it does a full system check. This only happens on the first start-up or after an improper power shut down. The system makes an audible “ticking” as it counts down through the initialization sequence. Please wait for the system to finish this process. If the system is restarted during the initialization sequence, the process starts over from the beginning. Please wait for the system to finish its initialization sequence to avoid start-up delays.

---

2. Move to a workstation with a web browser such as Microsoft Internet Explorer on the same subnet as PROVIGO 410. Type either http://192.168.0.1 or http://192.168.0.2 into the web browser address window to access the PROVIGO 410 RAIDconn.

Note: The workstation that will be used for system configuration should be on the same subnet as PROVIGO 410. Set the Workstation IP address to 192.168.0.x and the net mask to 255.255.255.0 before trying to access the PROVIGO 410 RAIDconn.

3. At the PROVIGO 410 RAIDconn login, enter the following login information when prompted.
   Password: root
   User Name: root

4. This will take you to Systems Properties page. You may use the Setup Wizard to guide you through the setup process and monitor your progress through the initial system setup.

Option 2 - Configuring your own IP address (advanced users)

If you do not wish to use the default IP addresses that ship with the PROVIGO 410, you will need to use the "Command Line Interface" (CLI) that is included with your system. If you have not done so already, consult with your network administrator for an IP addresses that you can use to set up PROVIGO 410 on your network.

Note: The CLI is for advanced users. It can be accessed via a terminal emulator running on a PC or laptop connected directly to the COM1 Serial port on the back of the PROVIGO 410. For more information on the CLI and a full list of commands. See "**Appendix B – Using the Command Line Interface**".)

1. First, connect the supplied serial cable to the COM 1 port on either controller.

2. Connect the opposite end of the cable to a workstation or laptop with a terminal emulation program (i.e. Microsoft Windows HyperTerminal).

3. You must use the following terminal emulation settings on the workstation/laptop to access the CLI:
   - Baud Rate – 115,200
   - Data Bits – 8
   - Parity - None
   - Stop Bits – 1
   - Flow Control Off

4. Power-on the enclosure, using the DC power/standby switch-on each power supply. The green DC power light on the power supply should turn on and stay on. The green activity light on each installed disk drive should be on. The PROVIGO 410 comes with its own pre-installed OS and RAID stack embedded in a Linux kernel. When the system is powered on, it boots automatically.

Note: When PROVIGO 410 is started for the first time, the system may take 1 to 2 minutes to initialize as it does a full system check. This only happens on the first start-up or after an improper power shut down. The system makes an audible "ticking" as it counts down through the initialization sequence. Please wait for the system to finish this process. If the system is restarted during the initialization sequence, the process starts over from the beginning. Please wait for the system to finish its initialization sequence to avoid system start-up delays.

---

5. Once connected and the PROVIGO 410 is running you will be prompted for following login information on the workstation/laptop terminal screen:
    - Login – root
    - Password – root

6. This will initiate the CLI network setup wizard.

Note: This wizard will only run automatically when you log in from the serial terminal as "root". It will NOT run when you login using telnet or log in as a non-root user. Once the network is configured, the wizard will not run automatically again. However, you can run the wizard manually by issuing the command 'dasetup' from the CLI (See **Appendix B – Using the Command Line Interface**).

---

7. Once launched, the wizard will walk your through the network configuration. The default values are in the brackets. You can enter the default value or type in the private network information you received from your network administrator.

8. Here is a sample of the what you will be asked to provide:

    - Enter the IP address for Controller 0 [192.168.0.1]:
    - Enter the IP address for Controller 1 [192.168.0.2]:
    - Enter the NetMask [255.255.255.0]:
    - Enter the Default Gateway [192.168.123.254]:
    - Enter the Name Server [209.116.241.10]:

9. Go to any workstation on the same subnet as the one you assigned to your PROVIGO 410 in Step 8 above and open the internet browser.

10. At the address line type in the IP address you assigned to the controller in Step 8.

11. Enter the following login information when prompted.
    - User Name: root
    - Password: root

This will take you to Systems Properties page. You may use the Setup Wizard to guide you through the setup process and monitor your progress through the initial system setup.
.

Introducing the PROVIGO 410 GUI

This section will review the major areas of the PROVIGO 410 GUI. You open the GUI by connecting to the IP address of a controller in the PROVIGO 410. This will open the login screen shown in **Figure 16**.

The default username on each controller is root. The default password is root.

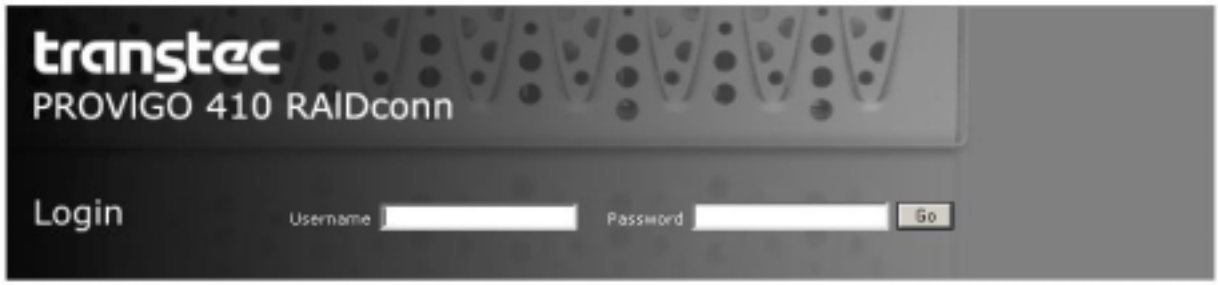

Figure 16 Login Screen

Once you have logged onto the controller, the System Properties page is opened. There are 6 primary parts of the GUI shown in **Figure 17**.

These are:

1. Banner area
2. Navigation Tabs
3. Setup Wizard
4. Navigation Tree
5. Work Area
6. Status Icon

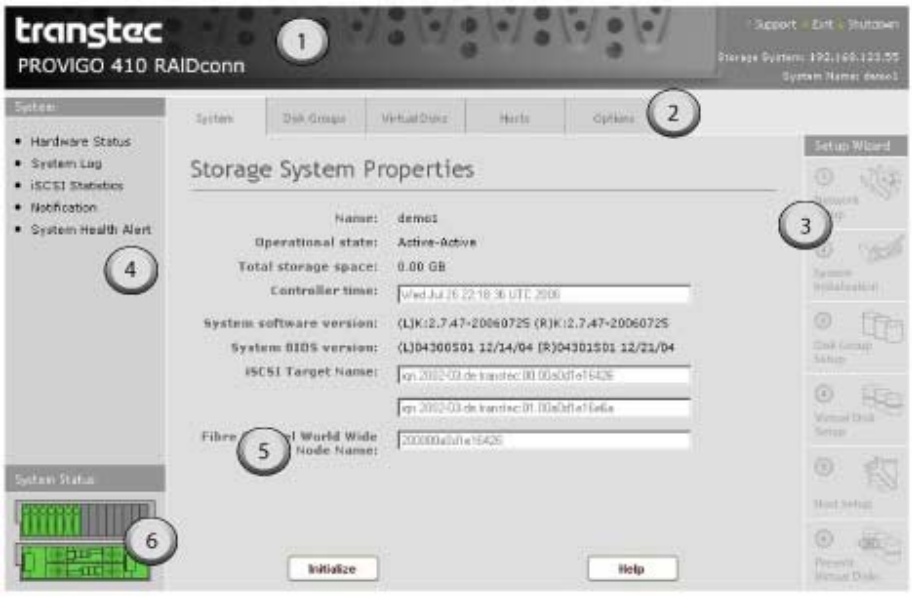

Figure 17 GUI Parts

In addition to the logo and product name, the Banner area also provides links to the support page, an exit link from the GUI, and a link to the Shutdown page. For easy reference, the Storage System IP address and the Storage System name are displayed as shown in **Figure 18**.

Figure 18 GUI Banner Area

The primary Navigation Tabs provide quick access to utilities that allow you to manage the PROVIGO 410 as shown in **Figure 19**. Each of these are discussed in greater detail in subsequent sections of this document.
The 6 tabs are:

1. System
2. Disk Groups
3. Virtual Disks
4. Hosts
5. Option

Figure 19 Navigation Tabs

The Setup Wizard shown in

Figure 20, provides a guide and progress bar that takes you through the six basic steps needed to configure and use your PROVIGO 410.

These steps include:

- Network Setup
- System Initialization
- Disk Group Setup
- Virtual Disk Setup
- Host Setup
- Present Virtual Disks

Figure 20 Setup Wizard

The Navigation tree shown in **Figure 21** provides a second level of navigation associated with the tab that is currently selected. For example, when System tab is selected, you can go to 4 additional sub categories.

These include:

Hardware Status
System Log
iSCSI Statistics
Notification

Figure 21 Navigation Tree

The Work Area contents, shown in **Figure 22**, vary based upon the Navigation tab you have selected. For each Navigation tab you will find information, input fields and action buttons associated with the specific Navigation tab or Navigation tree you have selected.

Figure 22 Work Areas

The System Status Icons shown in Figure 23 provides a quick visual representation of the hardware status in the PROVIGO 410. Healthy components are presented in green and faulty components are represented in Red. **Figure 23** shows how a drive color changes when a drive has a fault condition.

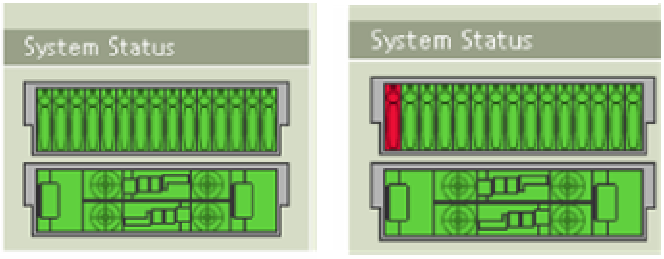

Figure 23 System Status Icons

Using the Setup Wizard

The Setup Wizard located on the right hand side of the GUI provides a guide and a progress bar for the six basic steps used to setup your PROVIGO 410 Storage array.

When you click on any of the steps in the Setup Wizard, that step is highlighted and the work area opens the associated pages for that step. Upon completion of a step, that area will be shaded to show that you completed that step in the setup procedure.

The shading of the progress in the Setup Wizard is reset when you log off from the controller.

Once you have completed the initial setup, you can use the Setup Wizard to jump to any of the steps you care to repeat. Those familiar with the system my chose to use the Navigation tabs to get to the same pages.

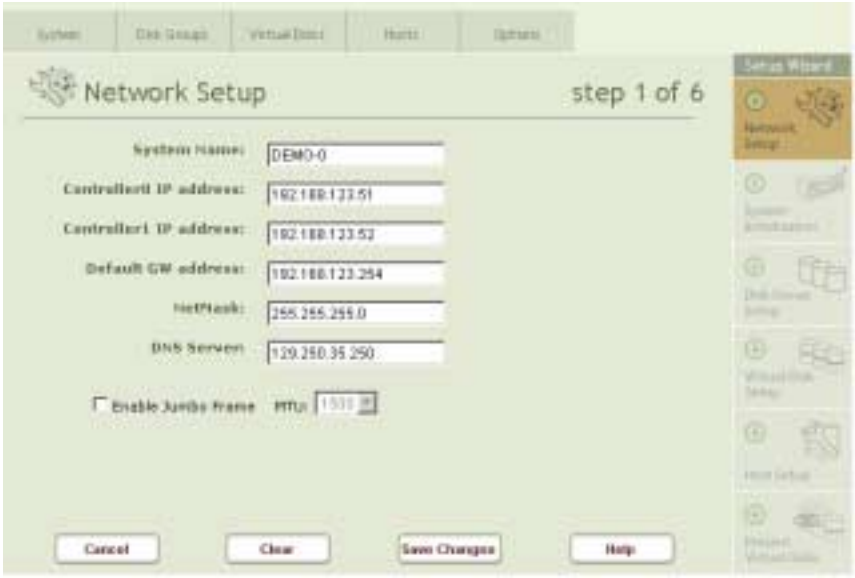

Figure 24 Step 1 - Network Setup

Step 1 – Network Setup

Use the Network Setup shown in **Figure 24** to configure you network connectivity to the system.

To do so:

1. Enter
   IP addresses for each controller
   Default Gateway (GW) address
   Netmask value
   Domain Name Server (DNS) IP address

2. Chose your Jumbo Frame setting
   Check the box and select a value for anything other than 1500 MTU

3. Click on Save Changes to have your setting take effect.

Note: If you change the IP address or netmask value from the address you used to log onto the controller, you will have to reconnect to the controller after clicking on Save Changes.

Step 2 – Initialize a Storage System

Step 2 System Initialization, shown in **Figure 25**, is used to define a system name, iSCSI target name, Fibre Channel World Wide Node Name and to set the date and time on the system. The work area has a slide bar to allow you to access the Fibre Channel World Wide Node Name setting. You may use the reset button to restore the default settings.

Once you have entered the system name, iSCSI target name, and Fibre Channel World Wide Node Name, you can click on Finish to save these setting or click on Next to set the date and time as shown in **Figure 26**.

Initialize a Storage System
step 2 of 6
Complete this step and click Finish to initialize your storage system in the simplest way possible. If you'd like more control over the initialization of your storage system, complete the step and click Next instead.
Enter a name for your storage system:
Cancel
Next
Finish
Reset

Figure 25 Step 2 - Initialize a Storage System

Initialize a Storage System
Continue with this step to initialize your storage system using advanced options. Click the Finish button to complete .
Use a custom date/time setting:
Date:
Time:
Time Zone:
PST
Cancel

Figure 26 Step 2 - System Initialization - Date & Time

Step 3 – Disk Group Setup

A Disk Group is a collection of disk drives that is created by selecting disks to become members of the disk group and building a RAID device with the selected disks. Clicking on Step 3, Disk Group Setup, shown in **Figure 27**, opens the Disk Group Properties page where you see an overview of the Disks in the system.

The fields displayed on the Disk Group Properties page are:

Operational State
: Shows the status of the disk groups

Total Disk Groups
: Shows the number of disk groups created

Total Grouped Disks
: Shows the number of disk drives that are members of a disk group

Total Spare Disks
: Shows the number of spare disk that are eligible to become members of a disk group through creation or as a hot spare to replace a disk that becomes faulty within a disk group.

Total Orphan Disks
: Shows the number of disks that have disk group information, but are not currently a member of a disk group

Total Not-Owner Disks
: Shows the number of disks that belong to a disk group owned by the remote controller. Disks are shown as Not-Owner disks when the remote controller is not available. This condition may occur when the remote controller is in the process of booting. You may wait for the remote controller to boot or use the Takeover DG to reassign these disks to the local controller.

Total Faulty Disks
: Shows the number of disks that have been declared faulty in the storage system.

Total Physical Disks
: Shows the number of disks installed in the storage system.

Buttons

Cancel
Return to the Storage Systems Properties page.

Create DG
Create a new Disk Group.

Discover DG
Manually discover Disk Groups.

Orphan to Spare
Convert Orphan disks into Spare Disks. This removes all Disk Group information from the selected disks.

Takeover DG
Reassign Not-Owner disks to the local controller

Help
Display the help page.

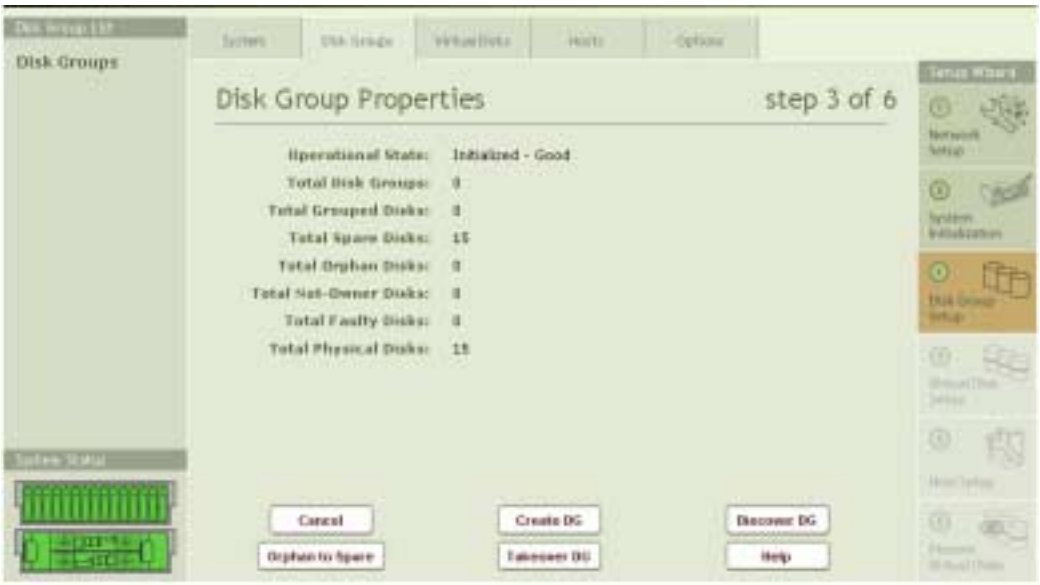

Figure 27 Step 3 - Disk Group Setup

To create a disk group, click on the Create DG button of the Disk Group Properties page. This will open the Create a Disk Group page shown in **Figure 28**.

Figure 28 Create Disk Group

Once you open the Create a Disk Group page, you need to enter a name for the Disk Group. The Disk Group name should not contain spaces or special characters. You will also select the drives to be used as members of the disk group. Then select a RAID level from the pull down menu as shown in **Figure 29**.

The system supports a maximum of 4 Disk Groups within the system. The RAID level selected for the Disk Group further imposes requirements associated with the RAID level and or the restriction on the maximum number of Disk Groups.

When creating a RAID 0 or RAID 5 Disk Group you must use a minimum of 3 drives. You may use a range of 3 drives to 15 drives when creating a RAID 0 or RAID 5 Disk Group. The same rule applies when expanding a RAID 0 or RAID 5 Disk group. You must use a minimum of 3 drives, and you can expand with a range of 3 drives to 12 drives.

When creating a RAID 1 or a RAID 10 Disk Group, you must use a minimum of 4 drives. You may create a RAID 1 or RAID 10 Disk Group with 4, 6, 8, 10, 12 or 14 drives. The same rule applies when expanding a RAID 1 or a RAID 10 Disk group. You must use a minimum of 4 drives and you may expand with 4, 6, 8, or 10 drives.

The following table describes the requirements for each RAID level.

Disk Group Parameters	RAID 0 & RAID 5	RAID 1 & RAID 10
Minimum number of drives in the Disk Group	3 drives	4 drives
Minimum number of drives for expansion	3 drives	4 drives
Maximum DG Capacity	16 TB*	16 TB*

Table 1 Disk Group - RAID Rules

- *Drive capacity continues to increase. Until drive capacities exceed 1TB per drive, it is not possible to have a 16 TB Disk Group in a single chassis.

Drives that are not assigned to a Disk Group are available as Global Spare unless they are in an Orphan state or Failed Disks. A Disk Group built as RAID 0 cannot use a global spare.

All drives within a RAID 1, RAID 5 or RAID 10 Disk Group should be the same size. The usable space on any given drive in the Disk Group will be equal to the size of the smallest drive in the Disk Group (except for RAID 0).

Figure 29 Create a Disk Group entry

You select a primary controller for the new Disk Group on the Create a Disk Group page. When using an iSCSI initiator, this setting will associate the Disk Group with the IP

address of the controller you select as defined in the network setup. Simply click on the radial button for the controller you wish to associate with the disk group.

Figure 30 shows the selection of Controller 0 for the new Disk Group. Clicking on the Finish button will create the Disk Group. Clicking on advanced will open the Advanced Features page shown in **Figure 31**, where you can set an Occupancy Warning level and add any comments to the Disk Group attributes.

Figure 30 Disk Group Controller Selections

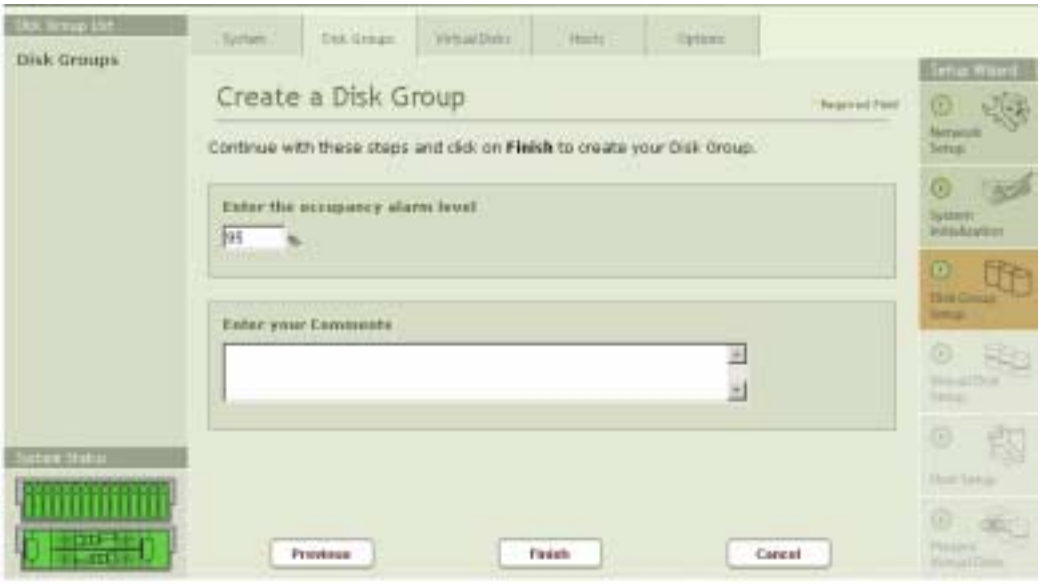

Figure 31 Create Disk Group - Advanced Features

Once you click on Finish, a progress page, shown in **Figure 32**, will open.

Figure 32 Create Disk Group - Progress Page

When the Disk Group has been created, a properties page for the newly created disk group is opened. You get to any specific Disk Group Properties page by selecting the

Disk Group you are interested in from the Disk Group List on the left hand side of the GUI. The currently selected Disk Group is highlighted in the Disk Group List as shown in **Figure 33**.

A RAID build is automatically started on the Disk Group when it is created according to the RAID level specified. The RAID is Building button is highlighted while the RAID is building. You can use the Disk Group while the RAID is building. You may also click on the RAID is Building button to open the RAID Build Progress page shown in **Figure 34**.

Figure 33 Disk Group Properties

The Disk Group Properties page for a specific Disk Group shows:

Name

The name of the Disk Group

Operational State

A Disk Group may be in one of the following states:

- Active is where the Disk Group is up and running with all member disks and no faults on any of the disks in the Disk Group.
- Inactive is where the Disk Group has been deactivated and is not running or available to any host.
- RAID Degraded is where the Disk Group is up and running but is either missing one or more Disk Group Members or where one or

more of the Disk Group members has been declared faulty. In this state, the Disk Group is usable but wit out RAID protection.

- RAID Unhealthy is where the Disk Group is missing two or more disk members. An Unhealthy Disk Group is not usable.

Total Disks

Shows the number of disks in the Disk Group.

Total Capacity

Shows the Disk Group capacity in Gigabytes.

Occupancy

Shows the percentage of the Disk Group that has been allocated to Virtual Disks.

Select Primary Controller

Shows the current Primary controller selection or allows you to change the primary controller selection through the pull down menu. Selecting an alternative primary controller will not take effect until you click on the Save Changes Button.

Protection Level

Shows the RAID protection level of the Disk Group. The available RAID levels are: RAID0, RAID1, RAID5, or RAID 10.

Disk Group Members

Shows the disk drives that are members of the Disk Group.

Comments

Shows any comments that were entered from the advance features of the Create Disk Group page or allows you to enter comments for the specific Disk Group. Comments will not be associated with the Disk Group until you click on the Save Changes button.

Buttons

RAID is Building

The RAID is Building button is shown while a RAID is building. This occurs during the initial creation of a Disk Group or when the RAID is being reconstructed. A RAID device is reconstructed when a member has been declared faulty and has been replaced by a global spare drive or by physically replacing the faulty drive.

Deactivate/Activate

Clicking on the Deactivate button will take a Disk Group off line, making all of the Virtual Disks unavailable to host on the network. For Disk Groups that have been deactivated, clicking on Active will make the Disk Group and all the Virtual Disks contained in the Disk Group available to hosts.

Expand

Click on Expand to increase the capacity of a Disk Group. You add members to the Disk Group in sets that are consistent with the RAID type.

Delete

Click on Delete to remove a Disk Group from the system. You must remove any Virtual Disk from the disk group before deleting the Disk Group.

Save Changes

Click on save changes to apply any changes to the Disk Group.

Help

Displays the Help page.

Disk Groups
Disk RAID Build in Progress
Type Percent Time(left) Speed Progress
resync 0.3% 219.2min 18548K/sec [>.....]
*Disk group can be used while RAID is building.
OK

Figure 34 RAID Build Progress

When the RAID build is done the RAID is Building button is no longer visible. **Figure 35** shows the Disk Group Properties page, as it will appear when the RAID build is completed.

Figure 35 Build Complete

Disk Groups may be moved between chassis. When disks belonging to a Disk Group are installed in a new chassis, they may appear as Orphaned Disks. The Discovery process, shown in Figure 36, will examine disks to determine if they belong to a Disk Group, if all the Disk Group members are present and check the state of the Disk Group.

You should de-activate a Disk Group before moving the disks to a new or different chassis. Removing multiple drives from an active Disk Group will create a failure scenario that makes the Disk Group unhealthy and unusable.

You may select Force Disk Group(s) online to bring a disk group online even when it is in a degraded state. You cannot use Force Disk Group(s) online to bring an unhealthy disk group online.

Disk Groups
• newdskgroup
Discovery Disk Group
Discovery examines the disks that are currently in the Orphan State. If the Disks in the Orphan State have the matching Superblock (the disks were in the same group), and all disks in the Disk Group are present and not faulty, Discovery will bring up this Disk Group. If you check the "Force Disk Group(s) online" option, Discovery will bring up the Disk Group even when it is in the degraded mode. For example, a RAID 5 Disk Group with a missing or faulty disk.
Select a Primary Controller
Controller 0
Controller 1
Force Disk Groups to be brought up in degraded mode
Force Disk Group(s) online
Cancel
Discover

Figure 36 Disk Group Discovery

An Orphan Disk is a disk that contains information about membership in a Disk Group that has either been removed from the Disk Group or is designated as an Orphan because there are not enough of the Disk Group members available in the system to build the Disk Group.

This condition can occur when a disk has been improperly removed from the system and a hot spare disk has replaced the Orphan Disk in the Disk Group. Since the Disk Group is rebuilt using the Spare Disk, the Orphan cannot be used in the Disk Group.

This condition may also occur when some disk from a Disk Group are moved to a new chassis without enough Disk Group members are present to reassemble the Disk Group. In this case, you can either locate and add the missing Disk Group members or convert the Orphan Disks to Spare Disks and rebuild the Disk Group from available Spare Disks. In this case all information on the disks will be lost.

Converting an Orphan Disk to Spare will remove all data from the disks. Only those disks that are in an Orphan state may be selected to be converted. Click on the Convert button in the Orphan disk to Spare disk page, shown in **Figure 37**, to convert the disks.

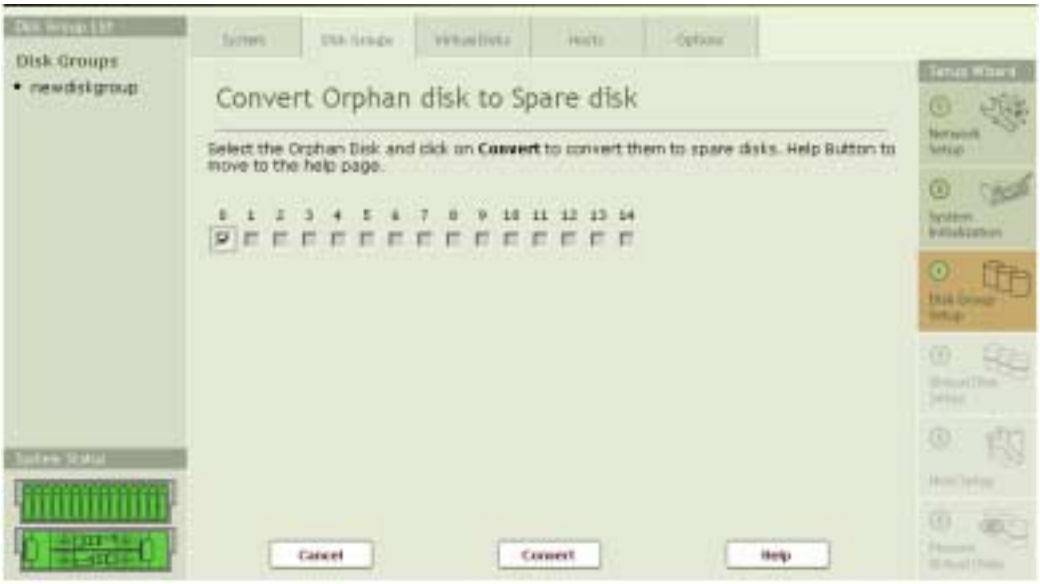
Disk Groups
• newdiskgroup
Convert Orphan disk to Spare disk
Select the Orphan Disk and click on Convert to convert them to spare disks. Help Button to move to the help page.
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14
Cancel
Convert
Help

Figure 37 Convert Orphan disk to Spare

Expanding a Disk Group

The ability to expand a Disk Group depends upon the number of drives and the type of RAID used in the disk group and having enough spare drives to complete the expansion. Expanding Disk Groups follows the same rules for each RAID type as when you created the Disk Group (*see **Table 1** on page **54** for additional information*).

Figure 38 Disk Group Properties – Expand

To expand a Disk Group, select the Disk Group you wish to expand from the Disk Group List on the left hand side of the GUI and click on the Expand button at the bottom of the properties page for that Disk Group shown in **Figure 38**.

This will open the Expand a Disk Group window shown in **Figure 39**, where you select the additional drives for the Disk Group.

Disk Groups
Expand a Disk Group
Select the disks to expand the Disk Group, and Click the Finish to expand the Disk Group.
Cancel
Finish
Help

Figure 39 Expand a Disk Group

Once you have selected the drives to add to the Disk Group, click on Finish to start the expansion. A progress window shown in **Figure 40** is displayed while the expansion is executed.

Figure 40 Expansion Progress

Step 4 – Virtual Disk Setup

Once you have built at least one Disk Group, clicking on Step 4, Virtual Disk Setup will open the Virtual Disk Properties pages show in **Figure 41**.

Note: The number of Virtual Disks (also known as logical volumes) in all the disk groups in the system must not exceed 256.

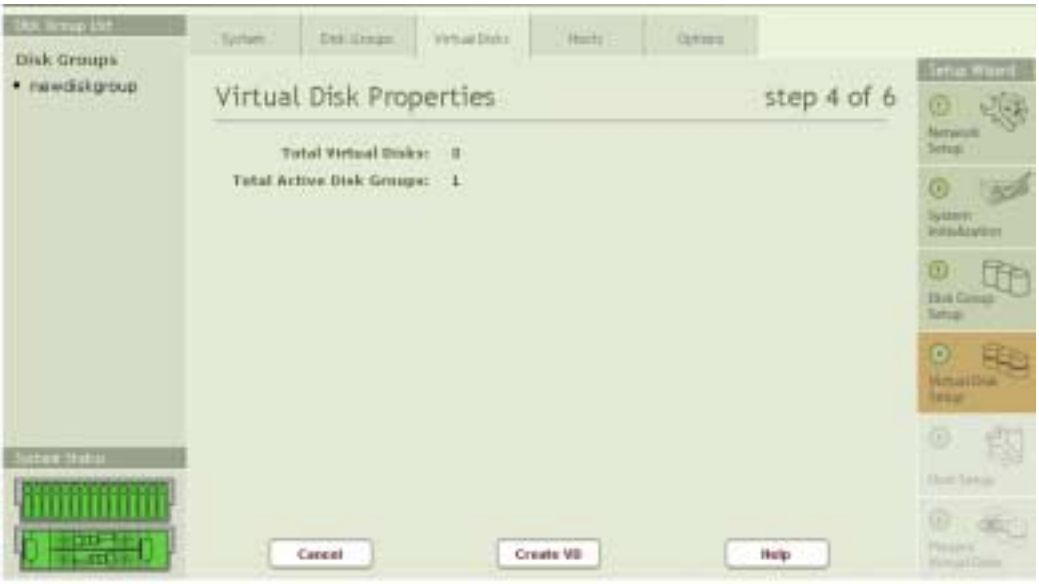

Figure 41 Virtual Disk Setup

The Virtual Disk Properties page in **Figure 41**, shows the Total number of Virtual Disk and Total active Disk Groups in the work area.

Buttons on the page include:

Cancel
: Returns the user to the System Properties Page.

Create VD
: Opens the page to create a Virtual Disk.

Help
: Displays help information.

Clicking on Create VD will open the page shown in **Figure 42**.

Figure 42 Initial Create VD page

To create a Virtual Disk, do the following:

1. Enter a name for the Virtual Disk.
2. Select the source Disk Group for the Virtual Disk as shown in **Figure 43**.

Figure 43 Create VD - Selecting Disk Group

3. Enter the size in GB from the available Free space in the Disk group

4. Select the cache policy for the Virtual Disk as shown in **Figure 44**.
   - Auto – provides mirrored cache and write back policies while two controllers are present and operation. Should one controller fail, or be removed the cache policy is switched to write through.
   - Write Through – passes data to the disk and cache simultaneously.
   - Write Back – does not provide mirrored cache. Write Back stores data in battery protected cache and writes to disks as cycles are available.

Figure 44 Create VD - Selecting Cache Policy

5. Select the read-ahead policy as shown in **Figure 45**.
   - Enable – allows read ahead operations. The read-ahead logic adjust for traffic patterns to improve bandwidth.
   - Disable – turns off read ahead operations. It may be desirable to disable read-ahead in environments that use very randomized small block read and write operations.

Figure 45 Create VD - Selecting read-ahead policy

6. Select VD Type as shown in **Figure 46**. SAN (iSCSI, FC) creates a LUN to be presented as an iSCSI or FC target.

Figure 46 Create VD - Selecting VD type

Figure 47 Virtual Disk Properties

The Virtual Disk Properties page shown in **Figure 47** displays:

Name
: The name of the VD. Works with the Save Changes button to modify this field.

VD Type
: SAN (iSCSI, FC)

Date/Time Created
: Displays the date and time the Virtual Disk was created

Operational State
- Active when available
- Inactive when the Disk Group is deactivated or doesn't have an owner

Disk Group
: The name of the Disk Group in which the Virtual Disk resides.

Cache Policy
: Auto, Write Through or Write Back. Works with the Save Changes button to modify this field.

Capacity
: The size of the Virtual Disk in GB

Read Ahead Policy
: Enabled or Disabled

Extend to new capacity
: Provides for entering a larger capacity. Works with the Save Changes button to take effect.

Buttons

Cancel
: Returns you to the Storage System Properties Page

Save Changes
: Works with Name, Cache Policy, Read Ahead policy and Extend Capacity field to make changes effective.

Presentation
: Used to open the Presentation window where you specify the host or hosts to connect to the LUN in SAN environments.

Take Snapshot

Opens the Snapshot window where you can take point in time read only snapshots.

Delete

Remove the Virtual Disk from the system.

Help

Displays the help page

Figure 48 Extending VD Capacity

Once you have created a Virtual Disk you can change several operational characteristics by entering a new value in a field on the Virtual Disk properties page and then clicking on the Save Changes Button. **Figure 48** shows entering a new capacity for the Virtual disk. This would be the size with the additional space added to the existing capacity. The additional space must be available in the Disk Group to allow the Virtual Disk to expand.

Manual, full copy point-in-time Snapshots (PIT) are supported. The Snapshots are read only and may be presented to a separate host from the original Virtual Disks. Because many iSCSI initiators cannot distinguish between the original Virtual Disk and the Snapshot, a Snapshot is only available to hosts other than the host where the source Virtual Disk is not presented. Read only Snapshots allow you to perform a backup operation as of the date and time the Snapshot is taken.

Snapshot Virtual Disks and the original Virtual Disk must reside in the same disk group. The combined Virtual Disk and Snapshot Virtual Disk size cannot exceed the Disk Group capacity.

You do not need to specify the snapshot size when creating a Snapshot Virtual Disk. The size of the original Virtual Disk is used to determine the size of the Snapshot Virtual Disk. The default initial size setting is 20% of the original Virtual Disk size.

You may set an initial size for the Snapshot based on the amount of data you expect in the original Virtual Disk. When data is sent to the source VD, data in the source VD is first copied onto the Snapshot before the new data is written into the source VD. It is possible to set the initial size of the Snapshot to a much smaller size than the original VD. If for example you set the size of the Snapshot to 20 percent, when more than 20 percent changes in the source VD, the Snapshot will double in size to 40 percent of the original. If data continued to change in the source VD, the Snapshot would continue to expand by increments equal to its original size until it reaches 100 percent of the source VD. If you expect that the contents of the original Virtual Disk will completely change, you can specify 100% as the size of the Snapshot Virtual Disk. This will prevent Snapshot Virtual Disk overflow.

It is important to understand that if the Snapshot needs to expand and there is no space left in the Disk Group to expand, the Snapshot will be disabled and deleted.

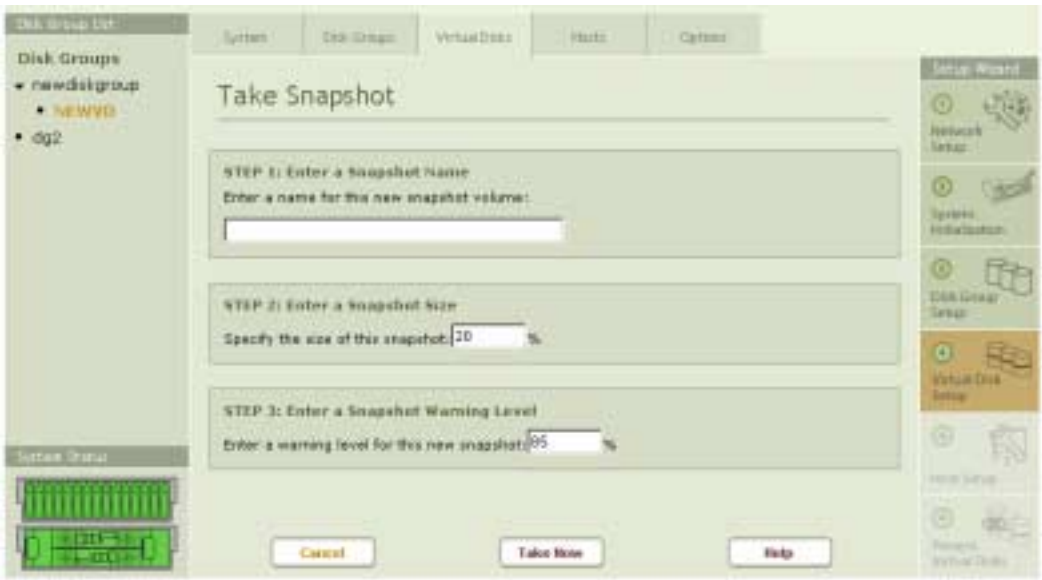
Disk Groups
newdiskgroup
NEWVD
dg2
Take Snapshot
STEP 1: Enter a Snapshot Name
Enter a name for the new snapshot volume:
STEP 2: Enter a Snapshot Size
Specify the size of this snapshot: 20 %
STEP 3: Enter a Snapshot Warning Level
Enter a warning level for this new snapshot: 95 %
Cancel
Take Now
Help

Figure 49 Virtual Disk Snapshot

Clicking on the Virtual Disk Properties page Take Snapshot button will open the Take Snapshot window shown in **Figure 49**. To take a Snapshot, enter a name for the Snapshot, select an initial size and warning level, then click on the Take Now button.

Figure 50 Snapshot Properties

When a snapshot is created the Snapshot Properties pages is opened. A Snapshot is shown in the Disk Group List under the source Virtual Disk. The Snapshot Properties pages shows:

Name
: The name of the Snapshot.

Date/Time Created
: The time the Snapshot was taken.

Operational State
: Read Only, as Snapshots may not be written to.

Disk Group
: The Disk Group the Snapshot resides in.

Virtual Size
: The size of the source VD and the maximum size of the Snapshot.

Capacity

The size of the Snapshot without expansion.

Percentage Full
The amount of the capacity that is in use.

Warning Level
The percentage at with a recipient on the notification list will be notified of usage and potential for expansion.

Buttons

Cancel
Cancel will take you to the Storage System Properties Page.

Save Changes
Works with the Name field to change the name of the Snapshot.

Presentation
Opens the presentation window to allow presentation to existing Hosts.

Delete
Removes the Snapshot from the system.

Help
Displays help information.

Figure 51 shows four snapshots taken where the first snapshot name is “snap” and snap1, snap2 and snap3 were created after snap. Selecting the snapshot on the left hand side of the GUI will display the snapshot properties page where you can see the Date/Time the snapshot was created.

Figure 51 Snapshot naming

Once four snapshots have been taken, any subsequent snapshots replace the oldest snapshot. Snapshots are listed with the oldest snapshot at the top of the list and the newest snapshot at the bottom of the list. The specific time a snapshot was taken can be determined by reviewing the Date/Time Created for any given snapshot in the Virtual Disk Snapshot Properties page.

Once you have created the Snapshot Virtual Disk, the Snapshot Virtual Disk is treated as a read only Virtual Disk that can be presented to an initiator as any other Virtual Disk.

When deleting Snapshot Virtual Disks, you must delete the oldest Snapshot Virtual Disk first. Snapshots Virtual Disks are listed in order from the oldest to the newest. Selecting the newest Snapshot Virtual Disk or an interim Snapshot Virtual Disk to delete will bring up an error message.

Expanding a Snapshot Virtual Disk

When application data is written into the original Virtual Disk, the old data will be copied into the snapshot Virtual Disk before it is overwritten by the new data. When the remaining space in a Snapshot Virtual Disk falls below 1GB, and if there is danger for space overflow, the Snapshot Virtual Disk will be extended if there is sufficient space in the disk group

The following guidelines apply to Snapshot Virtual Disk expansion.

1) Double the Snapshot Virtual Disk size up to 100% of the source if possible.
2) If there is not sufficient space available in the Disk Group, expand the Snapshot Virtual Disk as much as possible.
3) If there is no space available to expand in the Disk Group, recycle previous Snapshot Virtual Disk space by deleting the oldest Snapshot Virtual Disk and expand the newest Snapshot Virtual Disk.
4) If there is no space available to expand the Snapshot Virtual Disk and no previous Snapshot Virtual Disk to delete, abort the expansion.
5) If a Snapshot Virtual Disk becomes overflowed, it will be disabled and deleted.

Step 5 – Host Setup

Clicking on Step 5 in the Setup Wizards opens the Host Properties page that is used to describe a Host to the Storage array as shown in **Figure 24**.

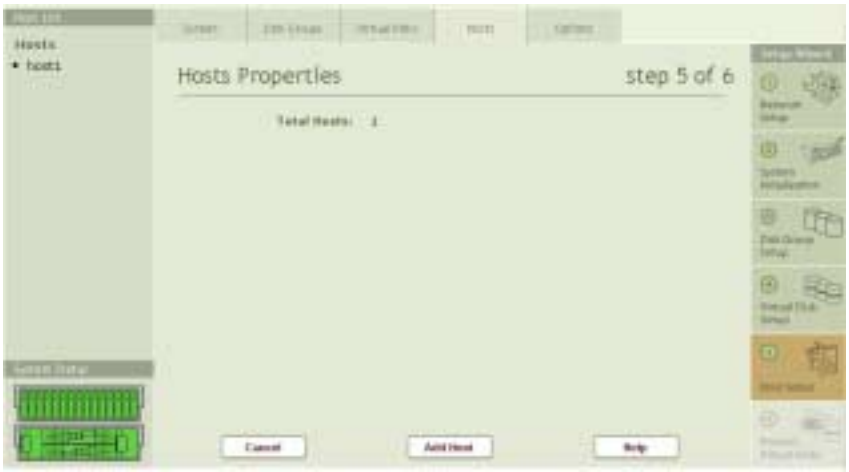

Figure 52 Host Properties Page

The left hand side of the screen shows any Host that are previously defined in the system. Clicking on the Add Host button will launch the Add Host wizard to guide you through the 4 simples steps required to describe a Host to the system.

Figure 53 shows the first 2 steps to add a host. Enter a Name for the new host in Step 1. Then in Step 2, select the interface to the host and enter either the name of the iSCSI initiator or the name of the Fibre Channel Port.

Figure 53 Adding a Host - Step 1 & 2

Figure 54 shows Steps 3 & 4 of adding a host. Use Step 3 to set up CHAP security. Enter the Target Secret and the Initiator Secret in the associated fields. Use Step 4 to add any comments about the Host that be displayed when you select the host from the Host List.

Click on Finish to complete the process of describing a Host to the system.

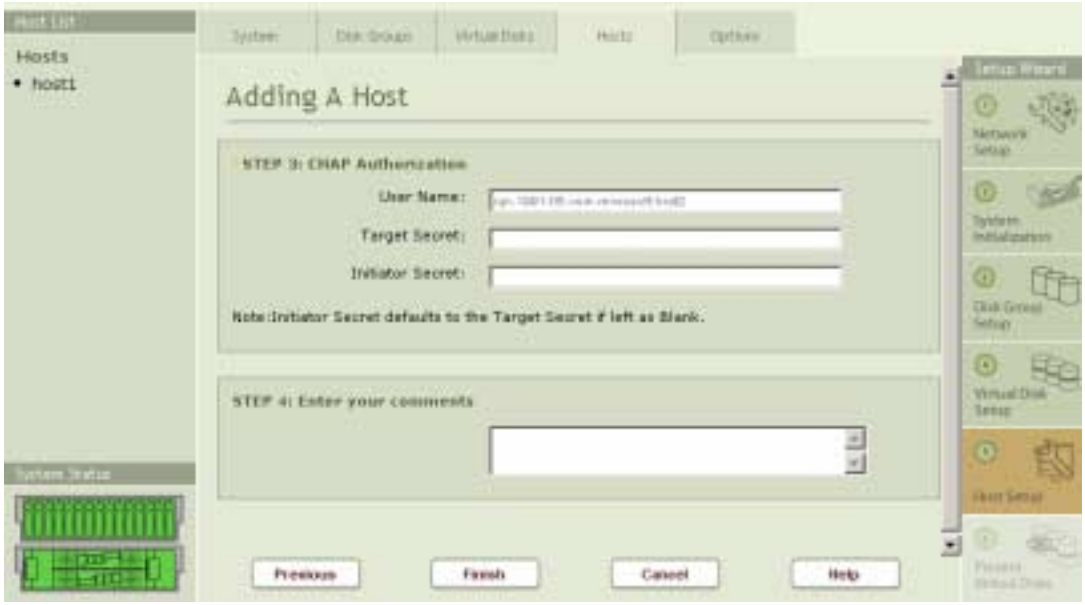

Figure 54 Adding a Host - Step 3 & 4

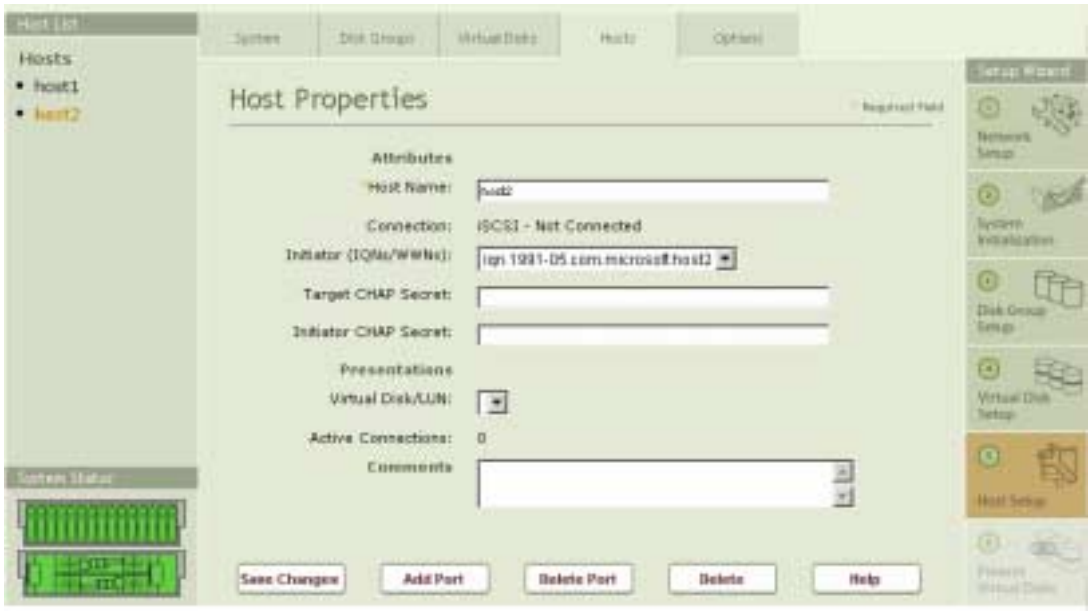

Figure 55 Host Properties

The Host Properties page for each host, shown in **Figure 55**, contains the following information about the host that is selected from the Host List:

Name
: Displays the current Host Name, or may be used with the Save Changes button to change the Host Name.

Connection
: Show the connection type and status

Initiator (IQNs/WWNs)
: Shows a list of iSCSI or Fibre channel ports that have been defined for the Host

Target CHAP Secret
: Shows the target CHAP secret or works with Save Changes button to add a target CHAP secret.

Initiator CHAP Secret
: Shows the Initiator CHAP secret or works with Save Changes button to add an initiator CHAP secret.

Presentations
: Shows a list of LUNS that have been presented to the host

Active Connections
: Shows the number of active connections to the Host

Comments
: Displays any comments that were added when the Host was defined or works with the Save Changes button to add comments to a Host.

Buttons

Save Changes
: Works with Name, Target CHAP Secret, Initiator Secret and comment fields to make changes take effect

Add Port
: Add more adaptor ports to the selected Host.

Delete Port
: Delete adaptor port from the selected Host. If the selected Host has only one port defined, the Host will be deleted.

Help

Displays the Help page.

Each host may support more than one adaptor. Click on the Add Port button to open the Adding A Host Port page shown in **Figure 56**. To add a port, enter the iSCSI Initiator name or Fibre Channel World Wide Port, select the interface and click on the Finish button.

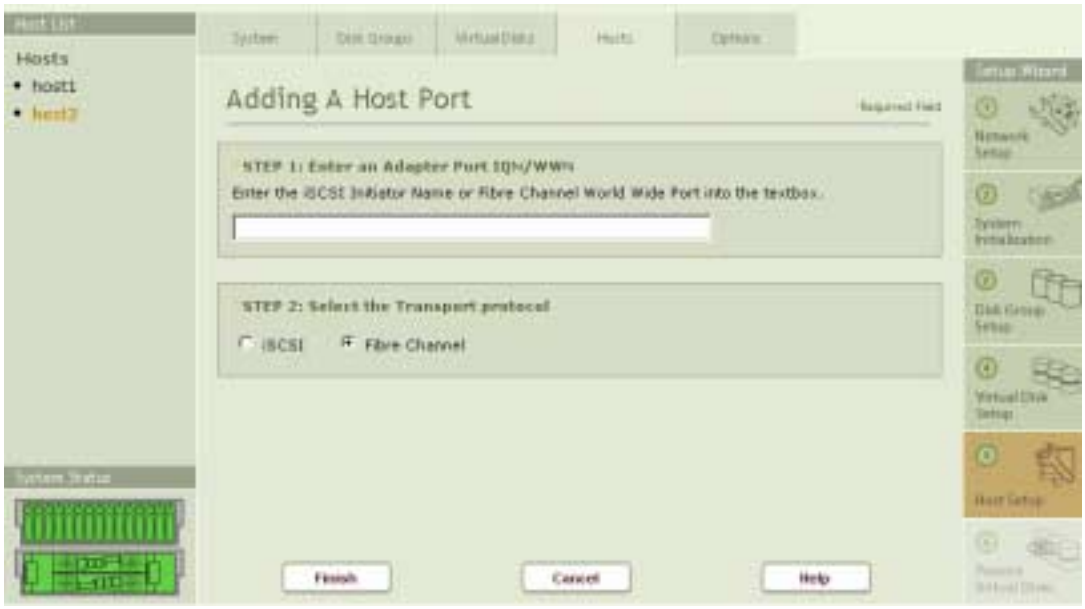

Figure 56 Adding A Host Port

Step 6 – Present Virtual Disks

The final step in setup is to present the Virtual Disk to Host. Clicking on Step 6 Present Virtual Disks opens the page shown in **Figure 57**. In this page, you select A VD to present from the pull down menu that shows available VDs and click on the Ok button.

Figure 57 Chose a VD to Present

The VD presentation page, shown in **Figure 58**, provides details about the presentations of the VD selected in the Disk Group List. Any host that the VD is currently present will show in the list box in the work area. Clicking on the Present button will open the Present Virtual Disk page shown in **Figure 59**.

Figure 58 VD Presentation

To Present a Virtual Disk, select the Host from the pull down list and click on the Finish button. If you wish to specify the LUN number to be associated with the Virtual Disk click on the Adv Option button to open the Present Virtual Disk advance feature window shown in **Figure 60**.

Figure 59 Present Virtual Disk Page

To change the LUN number associated with a Virtual Disk, select the desired LUN number from the pull down menu and click on the Finish button shown in **Figure 60**.

Figure 60 Present Virtual Disk Advanced Features

Configuring Options

Overview

The PROVIGO 410 allows you to perform additional task through the Options Tab. These include:

- Changing the system administrator password
- Changing the refresh rate to the browser
- Enable or disable System Health Alerts
- Enable or disable MPIO for the storage array
- Manage iSNS services
- Upgrade system software from the GUI

You can access the options by clicking on the Options tab of the system toolbar in the GUI. From the Options window you can select the area of interest by clicking on Go Now for the option you desire or by clicking on the item in the Option List on the left hand side of the GUI.

Figure 61 Options

Changing a storage system password

To change passwords on the system, click on the Change a storage system password icon in the options window to bring up the table shown in **Figure 62**. You may change the password for root or guest through this menu.

Figure 62 Change Storage System Password

To change the password:

1. Enter the user name
2. Enter the old password
3. Enter the new password
4. Re-enter the new password
5. Click on the update button

Changing the Display Refresh Options

By default, the GUI is refreshed every 5 seconds. You can change the refresh interval by clicking on the Display refresh options icon in the Options window. This will bring up the Display Refresh Options as shown in **Figure 63**.

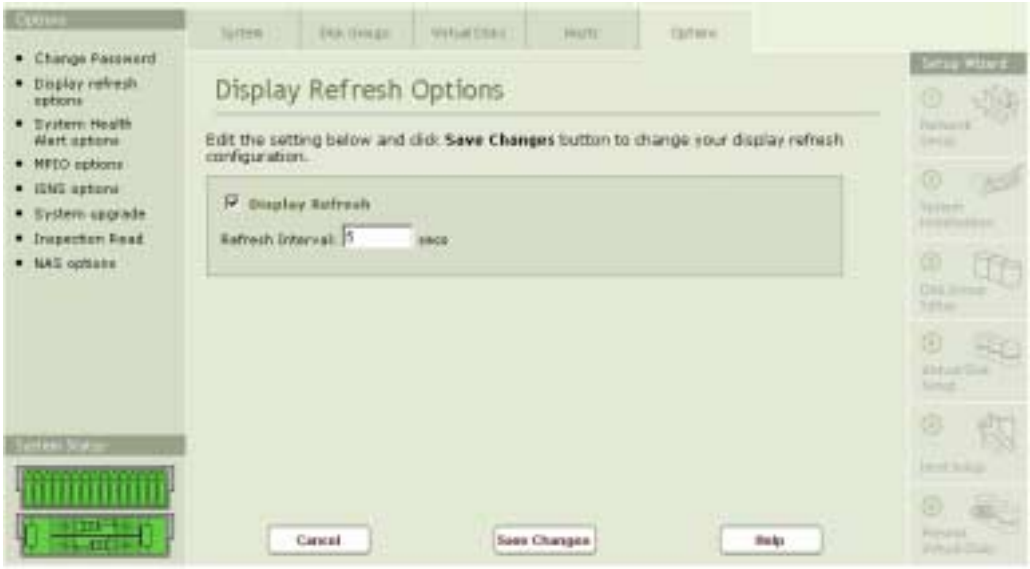

Figure 63 Display Refresh Options

To change the refresh interval:

1. Enter a new value in seconds
2. Click on the Save Changes button

Changing System Health Alert Options

System Health Alert Options may be enabled or disabled through the System Health Alter Options page. You can reset the alerts status by temporarily disabling then enabling the System Health Alerts. Click on the System Health Alert Option icon on the Options page to bring up the System Health Alert page shown in **Figure 64**.

Figure 64 System Health Alert Options

To change health alert options

1. Select Enable or Disable on the System Health Alert Options page
2. Click on the Save Changes button

Setting MPIO Options

The PROVIGO 410 support iSCSI-based MPIO with initiators that support MPIO such as Microsoft iSCSI initiator 2.0. To enable or disable the MPIO option, click on the MPIO Options icon in the Options page to bring up the MPIO Options page shown in .

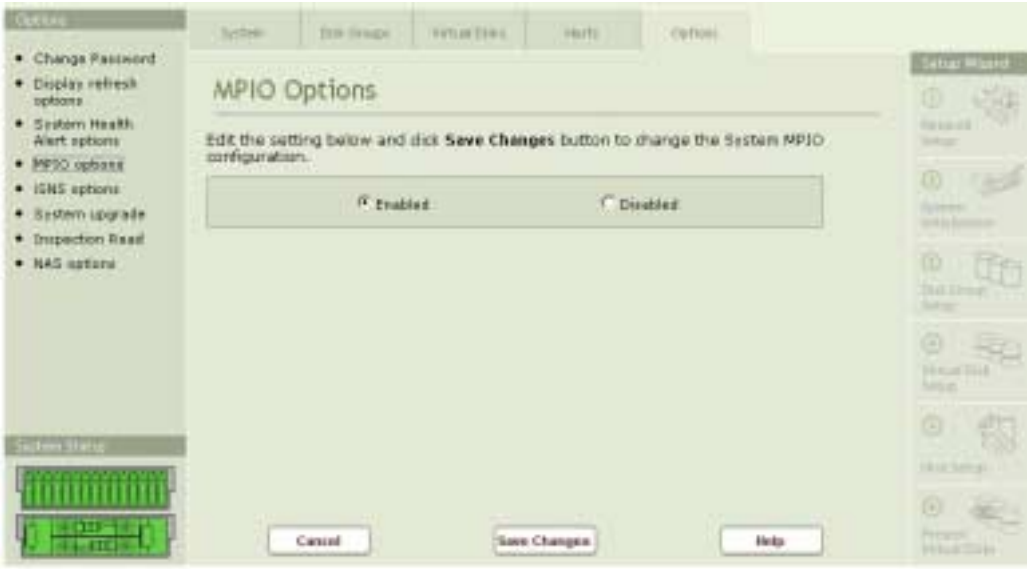

Figure 65 MPIO Options

To change the MPIO Option

1. Select Enable or Disable on the MPIO Options page
2. Click on the Save Changes button

Configuring the iSNS Client

The PROVIGO 410 provides an iSNS client for registration with an iSNS server. To fully utilize the iSNS service you must configure an iSNS server and configure your iSCSI initiator to use the iSNS server.

Once you have defined the IP address and alias for the iSNS Server, you must start the iSNS client on the target. Starting the iSNS client and saving changes will enable automatic iSNS client start up on any subsequent reboot.

To configure and use the iSNS target client, click on the iSNS Options icon on the Options page to bring up the iSNS Configuration page shown in **Figure 66**.

Figure 66 iSNS Configuration

To configure the iSNS client:

1. Click on the iSNS Options icon on the Options page
2. Enter the IP Address of the ISNS server
3. Enter the Server Alias of the iSNS Server
4. Select Start iSNS
5. Click on the Save Changes button

You can also determine the iSNS Server Status by Clicking on the Server Status button of the iSNS Configuration page. This will open a pop up window that will indicate rather the iSNS server you have defined is online or offline.

System Upgrade from the GUI

The PROVIGO 410 supports a system upgrade from the GUI after KDI (Kernel Downloadable Image) version 2.6.4. This feature allows you to select the new KDI from your local machine or another system on the same subnet as the system.

The upgrade process is initiated through the Upgrade Manager shown in **Figure 67**. Upgrade includes the following general steps.

- load a copy of the specified kdi into memory
- copy the kdi to the remote controller backup location
- reboot the remote controller with the backup copy moving IP address and disk groups to the local controller

- load the kdi into the backup location on the local controller
- reboot the local controller moving IP address and disk groups to the remote controller
- restore IP addresses
- restore disk groups to primary controller assignments (assuming minimum or no traffic)
- synchronize backup copy with new KDI

Messages are presented in the Upgrade Manger page Upgrade Status area showing the progress through the upgrade process. Command response and GUI refresh will slow during the final step that synchronizes the copies of the KDI.

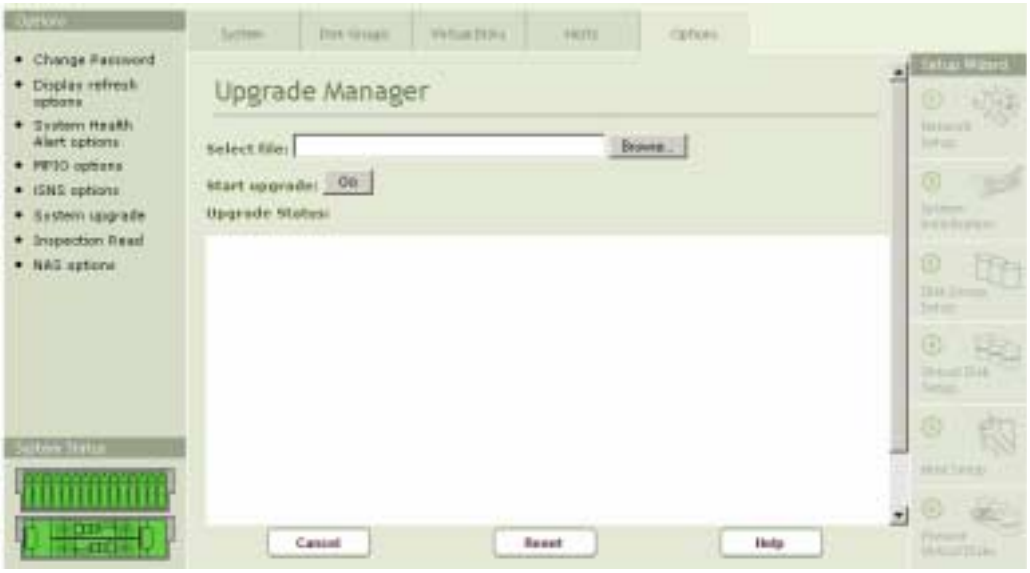

Figure 67 Upgrade Manager

To upgrade the system software:

1. Click on the System Upgrade icon in the Options page
2. Enter the software filename or use the Browse button to select a new software image
3. Start the upgrade by clicking on the GO button

Inspection Read

Inspection Read is an application designed to read from many of the disks in the storage array to proactively identify faulty disks before the user attempts to access any faulty regions.

Priority

Inspection Read may be run in the background at one of three different priority levels. Internally, this determines how much time is allocated to Inspection Read in the system. Inspection Read will automatically relinquish the CPU during heavy load on the system.

Logging

The Inspection Read process reports errors only on the local controller where the process was initiated. Errors and process output is recorded in /var/log/messages and KEM log files. Inspection Read creates a log message after each device has been successfully inspected.

Devices to Read

Inspection Read will access all sectors within each healthy, non-building, non RAID-0 Disk Group and spare disk drives.

Device Type	Action
RAID 0 Device	Do Not Read
Healthy RAID Device (not RAID 0, Not during initial build)	Read
Healthy RAID Device (not RAID 0, During initial build)	Do Not Read
Degraded RIAD Device (not rebuilding on spare disk)	Do Not Read
Degraded RAID Device (During rebuild on spare disk)	Do Not Read
Unhealthy RAID Device	Do Not Read
Orphan disk	Do Not Read
Spare disk	Read

Inspection Read monitors the status of Disk Groups and will stop an inspection if the status of the Disk Group changes to an unreadable state listed in the previous table.

Using Inspection Read

Inspection Read may be started from the GUI or from the CLI. This section describes using Inspection Read through the GUI. To learn how to use Inspection read from the command line interface, issues daisnpread –h.

Options GUI

To start Inspection Read from the GUI, select the Options Tab and click on Go Now under Inspection Read as shown in **Figure 68**.

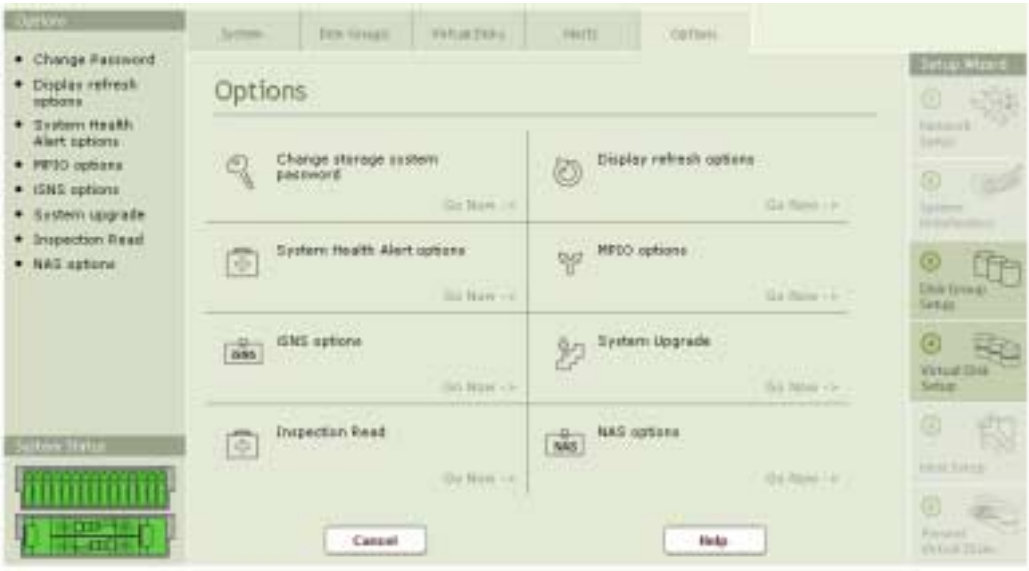

Figure 68 Options Tab

Inspection Read GUI Interface

To run Inspection Read from the GUI shown in **Figure 69**, you must select the priority to run the process at and any other options you desire then either click on the Start button or the Schedule button to start Inspection Read at a later time.

Once Inspection Read has started you can use the Stop and Resume buttons to pause or restart Inspection Read from the point it was stopped.

You can use the Progress button to see information about the progress of Inspection Read on the system as shown in **Figure 70**.

Checking the Repeat box will cause Inspection Read to run continuously until the user clicks on the Stop button.

Checking the Set Maximum Duration button allows the user to specify a maximum number of hours to run Inspection Read in the Maximum Duration field.

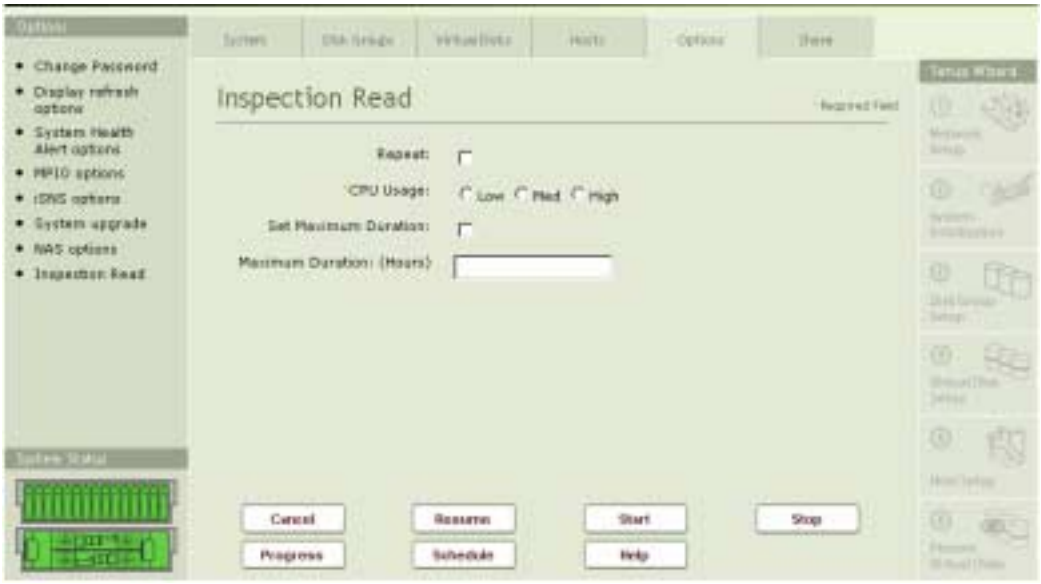

Figure 69 Inspection Read GUI Interface

Inspection Read Progress

Figure 70 shows the Inspection Read progress pages. The display fields include:

- Running – yes if Inspection Read is currently running
- Execution Time – show the amount of time that Inspection Read has been running
- Maximum Duration – Shows the total time allocated to run Inspection Read when Maximum duration has been set.
- Device list – Shows the devices and slots for each device under inspection. The Number of blocks inspected, total blocks to be inspected and percentage completed are listed for each device.

Change Password
Display refresh options
System Health Alert options
PHIO options
iSNS options
System upgrade
Inspection Read
NAS options
Inspection Read Progress
Running: Yes
Execution Time: 0 hours 0 minutes
Maximum Duration: 0 hours
RAID device md32: completed = 0MB, total = 128.375MB (0%)
RAID device md128: completed = 0MB, total = 0MB (0%)
Spare (slot 9): completed = 0MB , total = 239372.438MB (0%)
Spare (slot 10): completed = 0MB , total = 239372.438MB (0%)
Spare (slot 11): completed = 0MB , total = 239372.438MB (0%)
Spare (slot 12): completed = 0MB , total = 239372.438MB (0%)
Spare (slot 13): completed = 0MB , total = 239372.438MB (0%)
Back
Help

Figure 70 Inspection Read Progress

Configuring Notification

The PROVIGO 410 supports email notification for alerts, major and critical events. You specify email subscribers to be notified through the notification page of the GUI.

In order to receive email, you need to make sure you have correctly set up the default Gateway and DNS server in the network setup.

1. Select the Email Notification page from the ‘System Status’ menu as shown in **Figure 71**.

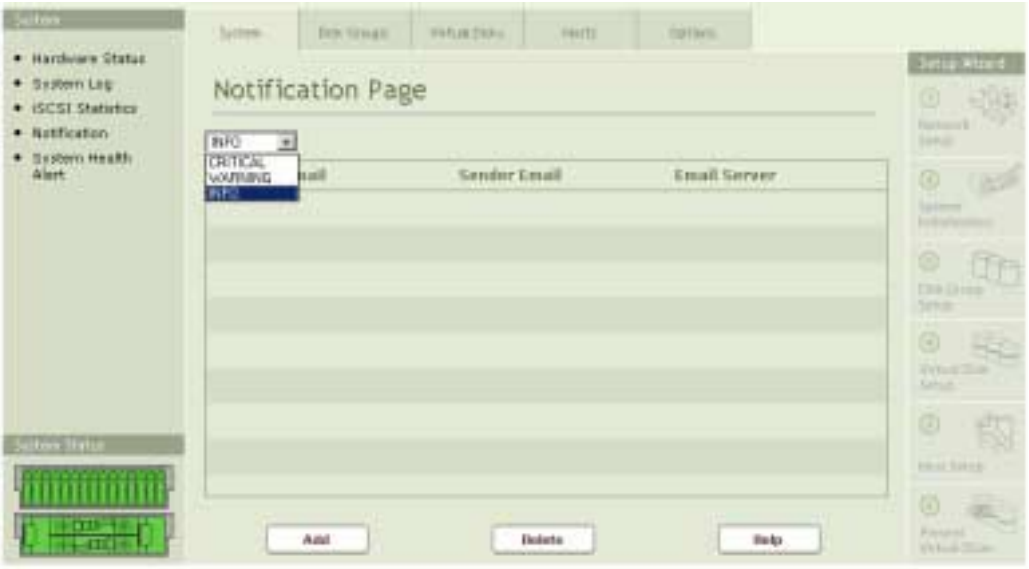

Figure 71 Email Notification

2. Click on “Add” to bring up the add user page as shown in **Figure 71**.

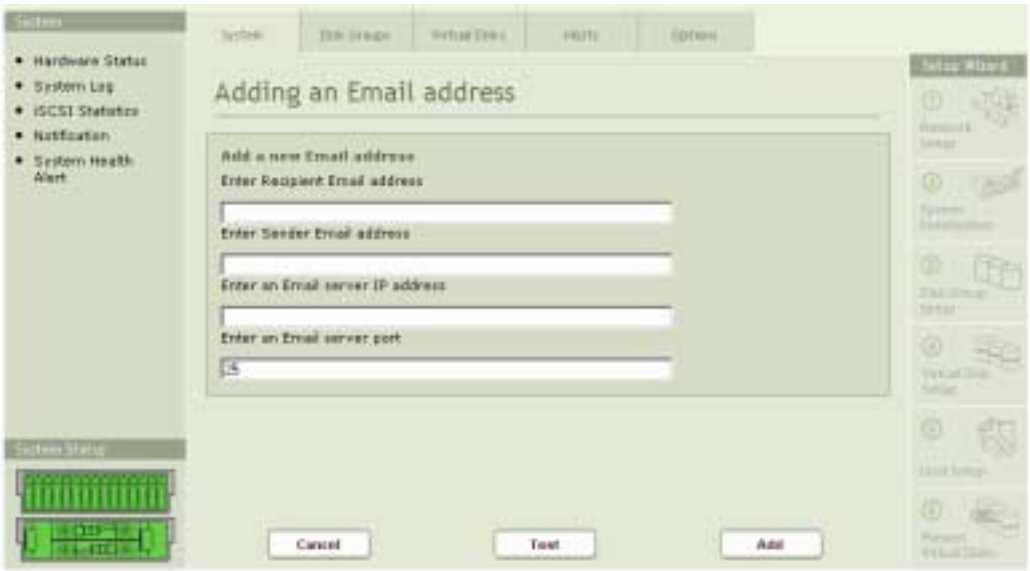

Figure 72 Adding Email Addresses

3. Enter
    a. Recipient Email Address
    b. Sender Email Address
    c. IP address of the mail server
    d. port number for the email service on the email server

4. click on the “Add” button as shown in **Figure 72**.

Using and Maintaining the System

This section includes instructions for using, maintaining and servicing your PROVIGO 410. It is organized from a hardware reference perspective and gives details on each hardware module and its associated configuration features. This section also reviews placement of LEDs and what they indicate.

System Overview

PROVIGO 410 chassis is a 3U/15 bay enclosure. It consists of a removable mid-plane and the metal housing that contains the other modules. The mid-plane provides connections for the disk drive modules, controller modules and power supply modules. The mid-plane is not hot swappable but can be easily removed for servicing.

The PROVIGO 410 chassis is in a rack mount form factor, which can be installed, in a standard Electrical Industries Association (EIA) 19-inch rack. Drives can be accessed from the front of the enclosure. All other field replaceable modules such as controllers, power supplies and fans can all be accessed from the back.

The PROVIGO 410 uses the same chassis as the PROVIGO 410. The controller and Hard Disk Modules have been redesigned to support SATA-2 disk drives and provide an additional SAS connection for capacity expansion on the controller modules.

Storage System Properties page

The System Properties Page is displayed when you log onto the system. This page provides an overview of the system as shown in **Figure 73**.

Figure 73 Storage System Properties Page

The fields in the Storage System Properties page are:

Name
: The hostname of the system.

Operational State
: Describes the operational state of each controller in the system.

Total Storage Space
: Displays the total capacity of the system in Gigabytes.

System Software Version
: Describes the software version of the local (L) and remote (R) controllers.

System BIOS version
: Describes the BIOS version of the local (L) and remote (R) controllers.

iSCSI target name
: Describes the target name for each controller.

Fibre Channel World Wide Node Name
: Displays the Fibre Channel World Wide Node Name.

Hardware Status

The hardware status page provides a visual representation of the system status. In addition moving the mouse over individual components provides additional details about the component. For example moving the mouse over a disk drive will display the manufacturer, model, size and firmware.

Figure 74 shows the Hardware Status for the PROVIGO 410. In either display, a failed component will display in red.

Figure 74 Hardware Status PROVIGO 410

System Log

The System Log are accessible through the GUI. To get to the System Log display, select System Log in the System list on the left hand side of the GUI. This will bring up the Log viewing screen as shown in **Figure 75**.

This screen allows you to view either the local or remote controller log entries. You can select the level of detail through the pull down menu as shown in **Figure 75**. Your choices include, info, warning or critical messages.

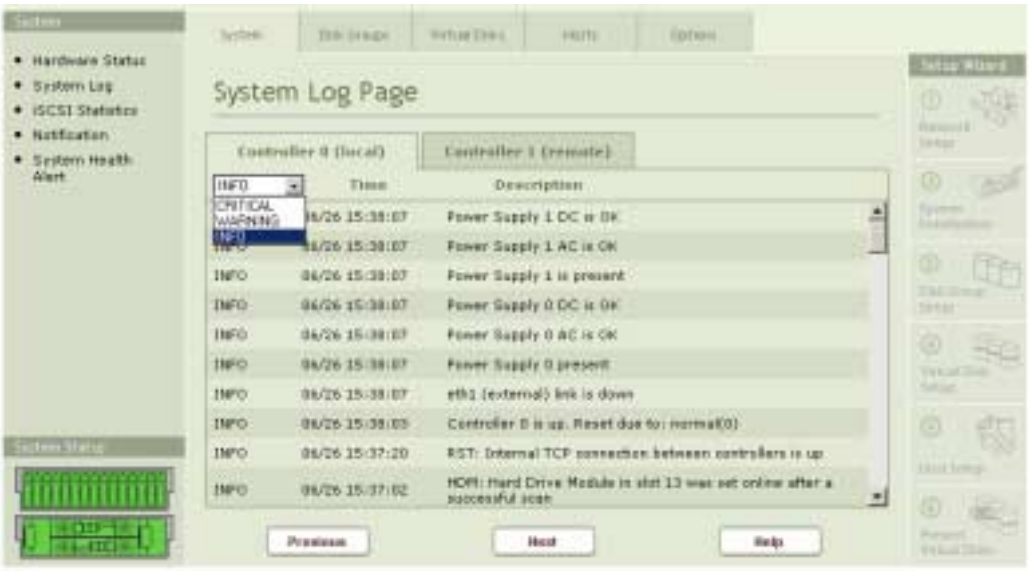

Figure 75 System Logs

ISCSI Statistics

Figure 76 shows the iSCSI statistics page. Each connection to the system is visible on the iSCSI statistics page. Information about the initiators, targets, connections, transmit and receive data are provided.

Selecting the Bandwidth button will open the bandwidth display shown in **Figure 77**. This display provides a graphical view of the total bandwidth to and from the chassis.

Figure 76 iSCSI Statistics Page

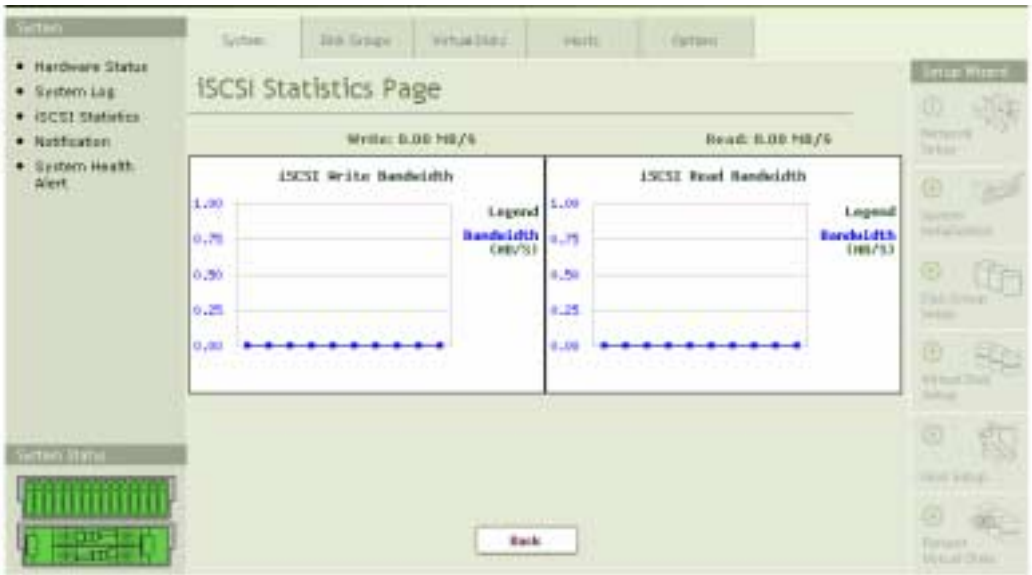

Figure 77 iSCSI Bandwidth Display

System Health Alert

transtec also equips PROVIGO 410 with a System Health Alert which is a flashing alert visible from the front of the system.

- **Description** – The System Health Alert lets administrators know that the system needs attention. The System Health Alert utilizes all the disk drive LEDs to flash a warning signal. All the LEDs on the disk drives blink approximately once per second. These lights continue to flash until redundancy is restored or the unhealthy condition is corrected. This visible alert is especially useful in dense rack implementations or situations where it is not easy to see or access the back of the system. Once alerted, you can log-on to the web-based PROVIGO 410 RAIDconn to ascertain the cause of the problem and use the system's Status Indicators (LEDS) to pinpoint the failed component and take corrective action.

- **Policy** – The System Health Alert is triggered whenever system availability degrades. For example, the Health Alert lights flash when component failure or removal causes system availability to drop below its original startup condition. It is also triggered by an unhealthy system state such as an exceeded temperature threshold.

> **Note:** PROVIGO 410 do not require redundant components to operate. If the system boots up without redundant power supplies, network links or controllers, the System Health Alert is NOT triggered.

To View the System Health Alerts, click on System Health Alerts in the navigation area of the System tab.

Figure 78 System Health Alert Page

Triggers and Corrective Action

This section describes specific events in the system that requires corrective actions and the corrective action associated with the event or alarm.

Temperature Threshold Exceeded

The system has multiple temperature sensors with factory thresholds. If any temperature sensor reading indicates an abnormal value beyond the threshold, the System Health Alert is triggered.

Corrective action:

- Log-on to PROVIGO 410 RAIDconn. Follow the menu tree on the left side of the screen to “System Status” and click on it. Check the graphical representation of the system for errors or click on the “System Health Alert” sub-menu for details on the alert.

- If the PROVIGO 410 RAIDconn indicates a faulty module, go to the system and look for a red Status Indicator. Replace the component (drive, fan, controller, power supply). Once the faulty module is replaced, the System Health Alert will stop flashing.

- Check for missing slot covers. Module slot covers are required to ensure proper system airflow. Replace missing slot covers to ensure temperature stays within specified system limits.

- If the PROVIGO 410 RAIDconn indicates a temperature fault and there are no faulty system modules, check for a sudden rise or drop in the room temperature. If the room temperature is beyond those listed in the system specifications section of this manual (See **Appendix C – System Specifications**) shut down the system until you are able to correct the problem.

Voltage Threshold Exceeded

The system is also equipped with voltage sensors. If any voltage sensor reading indicates an abnormal value beyond the factory threshold, the System Health Alert is triggered.

Corrective action:

- Check the Web-based console for the source of the problem.
- Check Status Indicator on each power supply. Replace if faulty.
- Ensure that the system is connected to an appropriate power source such as the rack's power distribution unit (PDU) or an uninterruptible power supply (UPS).

Fan Failure

Fan state and fan speed are also monitored. If a fan fails or if the fan speed falls below the factory threshold, the System Health Alert is triggered.

Corrective Action:

- Check the Web-based console for the source of the problem.
- Inspect the controller's "Tray 0" or "Tray 1" LED. If these are blinking or solid red replace the fans on that controller. (See **Controller Modules** on page **107**)

Degraded System Availability

When the PROVIGO 410 is turned on, it captures an inventory of installed components. Once the user configures disk groups, the system also tracks the RAID configuration information. If at any time, the system falls to lower system availability than at startup, the System Health Alert is triggered. This can happen if a redundant component fails or a RAID 1, 5 or 1+0 disk fails.

Corrective action:

- Check the PROVIGO 410 RAIDconn for the source of the problem.
- Restore redundancy by replacing the failed controller, power supply or failed drive. Once system availability is restored the System Health Alert will stop flashing.
- If the system has a global spare, the System Health Alert will stop flashing once the RAID group is automatically rebuilt by the system.

Unhealthy RAID Disk Group

If a RAID Group experiences a drive fault that is unrecoverable, the System Health Alert is triggered. For example, this can happen if a RAID 0 disk group experiences a drive failure or a RAID 5 disk group experiences two consecutive drive failures.

Corrective action.

- Recover the disk group from a recent back-up.

ECC Memory Error

The system uses ECC memory on each controller. If an ECC memory error is detected by the system, the System Health Alert is triggered.

Corrective action:

- Check the Web-based console for the source of the problem.
- Replace the cache memory module if necessary.

Battery Fault (Cache Battery)

Each controller comes with a battery backup unit to protect the cache. If the battery fails or stops holding a charge, the System Health Alert is triggered.

Corrective action.

- Check the Web-based console for the source of the problem.
- Replace the controller if necessary.

Faulty Battery Charger

The backup battery subsystem includes a charger for the backup battery. If the charger becomes defective, the System Health Alert is triggered.

Corrective action.

- Check the Web-based console for the source of the problem.
- Replace the controller.

Network Connection Lost

If a network link is lost, the System Health Alert is triggered.

Corrective action:

- Check the Web-based console for the source of the problem.
- If there are no system problems, check your network devices or initiators for the source of the problem.

Turning the System Health Alert Off

Once a system administrator identifies the problem and schedules corrective action, he may want to turn off the System Health Alert. To shut off the flashing lights, the administrator needs to:

- Log-on to the PROVIGO 410 RAIDconn.
- Once logged-in, click the “Options” menu at the upper left portion of the main screen.
- In the options menu, click “System Health Alert Options”
- There are two option buttons “enable and “disable”. Hit the disable button to shut off the System Health Alert.

Note: Although the System Health Alert is disabled, the system will continue to log error notifications in PROVIGO 410’s system log. Administrators can view the notifications in PROVIGO 410 RAIDconn’s “System Log” or “System Health Alert” sub-menus.

You may find it useful to disable and then immediately enable System Health Alerts when you have identified a problem with Power Supplies, Cables or a Controller. This will establish a new base line for the system configuration and ensure you are notified of any new alerts that may occur.

System Health Alert Error Notifications:

- The other controller is down
- RAM ECC error is detected

- Battery not present or not in normal charging state
- Battery is discharging unexpectedly
- The charger is bad
- Temp sensor reading T1 of xx.xx C > alarm threshold 65 C
- Temp sensor reading T2 of xx.xx C > alarm threshold 55 C
- Temp sensor reading T3 of xx.xx C > alarm threshold 55 C
- Voltage sensor reading V4 of xx.xx out of range [2.37, 2.62]
- Voltage sensor reading V4 of xx.xx out of range [1.35, 1.60]
- Voltage sensor reading V4 of xx.xx out of range [3.00, 3.60]
- Voltage sensor reading V4 of xx.xx out of range [4.60, 5.25]
- Voltage sensor reading V4 of xx.xx out of range [11.40, 12.60]
- Fan tray 0 - Fan 0 is missing or down
- Fan tray 0 - Fan 1 is missing or down
- Fan tray 1 - Fan 0 is missing or down
- Fan tray 1 - Fan 1 is missing or down
- A disk group becomes degraded or unhealthy
- Power supply 0 is not operational
- Power supply 1 is not operational
- Ethernet link on eth0 is down
- Ethernet link on eth1 is down
- Ethernet link on eth2 is down
- Ethernet link on eth3 is down
- Incompatible KDI – The other controller is upgrading me
- Incompatible KDI – The other controller is syncing itself
- Disk Group ([name]) ([raid 1/10]) has x degraded raid devices, cannot be discovered if rebooted

Hot Swap Module Overview

Each PROVIGO 410 comes with multiple disk drive modules and redundant Fan modules. Therefore, the Disk drive modules and fan modules may be hot-swapped in most systems.

Some systems may have only one Controller module or one power supply module. The Controller module and Power supply modules are can be hot swapped if an optional Controller module and Power supply module are present.

Caution:

- Do not remove an active disk from a RAID0 Disk Group, or an active disk from a degraded RAID1, RAID5 or RAID1+0 Disk Group
- Do not attempt to remove the mid-plane while the system is

powered on. Power must be off and all components must be uninstalled before the mid-plane can be removed.

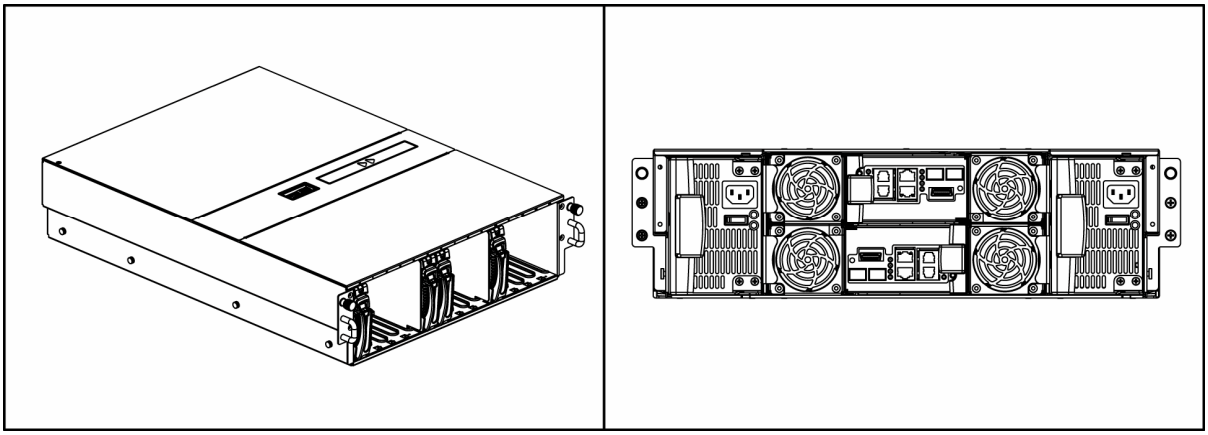

Figure 79 PROVIGO 410 Front and Rear View

Controller Modules

The controller modules are designed to be hot plug able and they fit into one of the two controller slots of the PROVIGO 410. The controller modules monitor and control every aspect of the storage subsystem. They control chassis management, control path and storage data path.

Controller Features

- Pentium CPU.
- Mirrored copy of the embedded RAID kernel
- Two PCI-X buses
- Two RJ45 Gigabit Ethernet ports
- Two DDR DIMM slots that support up to 2GB of memory running in 144 bit interleave mode. (The system ships with 1 GB)
- One Lithium-ion battery module to protect the write cache against any power failure for 72 hours.
- Two hot plug able fan modules per controller.
- Four Serial ATA controllers, 15 Serial-ATA channels.
- Two serial console ports. COM1 for Serial Terminal Connection and COM2 for UPS management.

Controller Features

The controller module has various status indicator lights (See **Figure 80**). Refer to table below for specific controller module conditions and corresponding LED color combinations.

①	CTL FAIL	Solid Red - Controller Failure
②	TRAY0	Blinking Red – One fan in Tray0 not working.
		Solid Red –Both fans in Tray0 not working
③	TRAY1	Blinking Red – One fan in Tray1 not working.
		Solid Red - Both fans in Tray1 not working
④	HB	Blinking Green – Controller operating normally (heartbeat normal).
⑤	Ethernet Link (ETH0 & ETH1)	Solid Green - Ethernet connection exists on this port. Blinking Amber – Ethernet activity on this port.
⑥	FC Link (FC1 & FC2)	All Off – Power Off Green On – Online Amber On – Link Established Amber Flashing – Loss of Synchronization Amber Flashing – Locator Beacon Green & Amber Flashing - Error

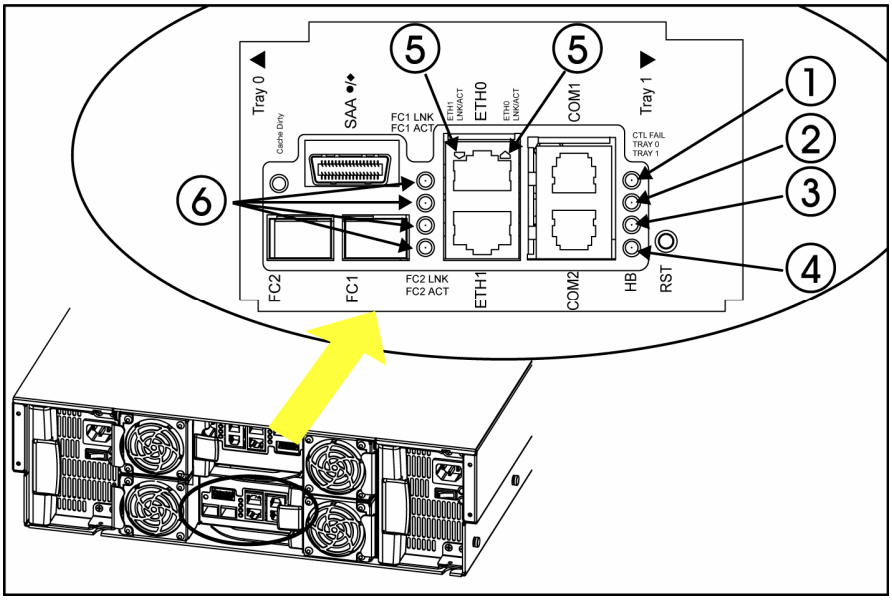

Figure 80 PROVIGO 410 Controller Component Locations

Controller States

PROVIGO 410 has three controller states. These are not configurable. The state is dependent on the purchased options, the physical configuration of the system as well as the functional state of the controllers. The following are the three possible controller operating states:

Active/Active Mode

With two controllers installed and auto cache enabled, the controllers function in active-active mode to improve both performance and data availability. Each disk group is assigned a "primary" controller to take advantage of a direct route to the controller. So, optimal performance is achieved when traffic goes through both controllers to their associated disk groups. For availability, both controllers can serve any disk groups within the system. When there is a controller failure, the IP address and disk groups assigned to the failed controller will automatically "migrate" to the surviving controller without user intervention.

Active/Takeover Mode

In a system with dual controllers installed, if one controller fails or resets, the surviving controller under PROVIGO 410's automatic fail-over mechanism will take over its IP address, target name and disk groups. This is called Active/Takeover Mode. The surviving controller will act as two controllers, one real and one virtual. This feature allows PROVIGO 410 to continue to operate without interruption in the event of a controller malfunction or failure. When in Active/Takeover Mode, the "Storage System

Properties" tab of the PROVIGO 410 RAIDconn will show that the system is running in this mode (See **Figure 82**).

When you log onto a system in Active-Takeover mode, a warning message appears as shown in **Figure 81**.

Your current Controller is in the Active-Takeover state. The remote controller has already come up. You can click on the System tab, then click on the [Ctrl Giveback] button in the System Properties page to give the ip address and target name back to the remote controller.
OK

Figure 81 Active-Takeover warning

To reset the system to Active/Active mode the user will need to manually reassign (fail-back) the IP address and disk groups to the new controller. (See **"Manual "Giveback" of IP Address**" on page **111**, and **"Fail-back" of Disk Groups**" on page **112**). Fail back must be manual to minimize the risk of data loss due to a hot-swapped controller incorrectly changing the IP address or bringing disk groups online that have already been brought online by the other controller.

Figure 82 Active/Takeover Mode

Active/- Mode

In a system with one controller installed, the system runs in Active/- mode. This means all disk are assigned to one controller. If the controller fails, a replacement must be installed and the system rebooted to reestablish availability of the system. A hot swapped controller will automatically reboot the system and regain access to the disk groups. This means that some downtime will be experienced in the event of a controller failure.

Manual "Giveback" of IP Address

When the failed controller comes back after reboot or is replaced with a new controller, it will detect that the remote controller is in the Active/Takeover mode, it will not attempt to bring up its external IP interface. When the user is ready to bring the system back to the Active-Active mode, they can log-on to the PROVIGO 410 RAIDconn and click on the ctrl Giveback button to move the IP address back to the original configuration.

Once in PROVIGO 410 RAIDconn, the user needs to click the "Storage System" text at the top of the menu tree on the left of the screen (See **Figure 83**). This will show the "System Properties" screen. A new button, "Ctrl Giveback", becomes visible at the bottom of the "System Properties Window" in the Active/Takeover Mode when a replacement controller is installed or the failed controller is reset. Press this button and the system will giveback the target name and IP address to the new controller.

Figure 83 IP Address Giveback

"Fail-back" of Disk Groups

Fail-back of Disk Groups must be manual so that Disk Groups that have been reassigned to a surviving controller, and may be active, are not taken offline due to automatic reassignment to a hot swapped controller. You must manually reassign disk groups back to the original controller. The ability to reassign a disk group from one controller to another without system interruption is a standard feature of PROVIGO 410.

The reassignment of a disk group to a different controller requires you to log-on to PROVIGO 410 RAIDconn.

Once in PROVIGO 410 RAIDconn, the user needs to click the desired disk group shown in the menu tree on the left side of the screen (See **Figure 84**). In the "Disk Group Properties" window, go down to the "Select Primary Controller" box to select the newly replaced or reset controller. Then select, "Save Change" button. Do this for each Disk Group you want to reassign to the new controller.

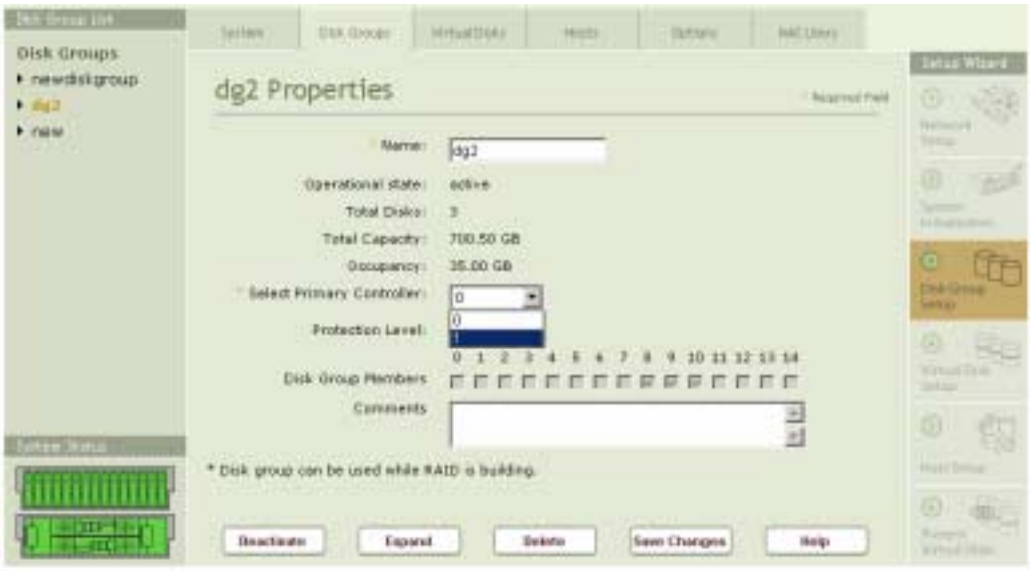

Figure 84 Disk Group Fail-back

Replacing a Failed Controller (Redundant System)

Note: Unless you have a particular reason to do so, do not remove power from the enclosure to complete this procedure.

Note: Take care to follow any special instructions shipped with a replacement controller module.

Attention: A controller module is sensitive to electrostatic-discharge (ESD). Use care in handling the controller and its components. (Refer to "ELECTROSTATIC DISCHARGE" on page 12for additional handling instructions)

1. To start the replacement process, first disconnect the controller's host interface cables and all other cables (if present) from the failed controller module.

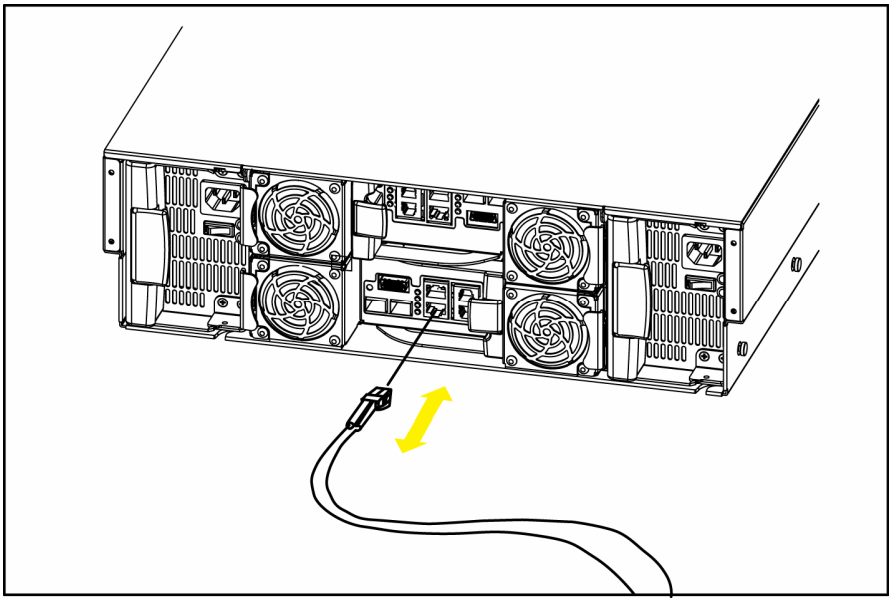

Figure 85 Disconnecting Connections from PROVIGO 410 Controller Module

2. Pinch the locking mechanism ① and swing the lever to the right. Open the lever fully. This action unplugs the controller module from the system enclosure.

Note: Do not remove any power cables when hot swapping controller modules.

3. Pull the controller module out of the enclosure, supporting it from below (See **Figure 86**).

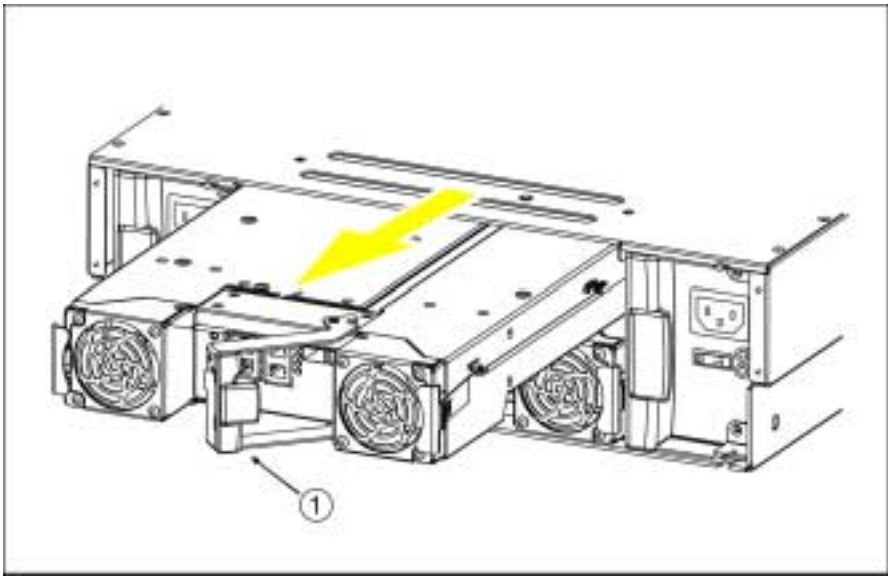

Figure 86 Removing the Controller Module

4. Replace parts in the reverse sequence.

5. Once the new controller is installed, the system will automatically recognize the controller and initialize it so that it is available to the system. This happens without any interruption to the system. The system is now fully redundant and able to recover from a controller failure.

Note: When a failed controller is replaced, the storage is not automatically reassigned to the new controller. Reassignment must be done manually.

6. The reassignment of a disk group to a different controller requires you log-on to PROVIGO 410 RAIDconn. In PROVIGO 410 RAIDconn, you click the desired disk group shown in the menu tree on the left side of the screen (See **Figure 84**). In the “Disk Group Properties” window, go to the “Select Primary Controller” box to select the newly replaced or reset controller. Then select, “Save Change” button. Do this for each Disk Group you want to reassign to the new controller. If the new controller is not listed, pull and reinsert the new controller and repeat each step.

Caution: Do not insert hands or tools into empty slots

Replacing a Failed Controller (Non-redundant System)

1. Stop all operations to and from all host systems that are trying to access this controller module.
2. Exchange the controller module for a new one. Follow the instructions in the previous section, “Replacing a Failed Controller (Redundant System)” to exchange the controller module with a new one.
3. Reconfigure the new controller module if necessary
4. Restart and/or resume operations from all the host systems.

Upgrading from a Single to Dual Controller

Attention: Do not remove power from the enclosure to complete this procedure.

1. Remove the slot cover from controller module slot.
2. Install the new controller module into the open controller slot.
3. Install host interface cables.
4. The system will automatically initiate the controller.
5. If you setup two IP Addresses when asked in your initial configuration of your system, the IP addresses will automatically be assigned to the new controller.

6. If you did not assign a new IP Address or if you wish to add storage or re-provision the established disk groups, log-in to the PROVIGO 410 RAIDconn.

Connectors

This section shows the locations of the external connectors of the PROVIGO 410. All cable connections are made at the rear of the enclosure (See Figure 87).

1. **Serial port connectors (RS232 port) -** COM1 is for serial terminal connection. COM 2 is for a UPS connection.

2. **Ethernet connectors -** PROVIGO 410 provides two Gigabit Ethernet ports. These two ports are set in a “Teaming” configuration from the factory.

3. **Fibre Channel connectors (available only on FC models)** - Two ports are provided per controller for Fibre Channel host connectivity.

4. **Power Connector** - Provides AC power connection. There is a power connector on each power supply module. The power connector connects a PROVIGO 410 to the AC power supply in the rack. The connector is a 10-amp, three-pin, polarized, IEC 320 input connector.

5. **SAS Connector** – Provides a SAS interface for expansion chassis.

Attention: Use only the power cables that are supplied with the enclosure.

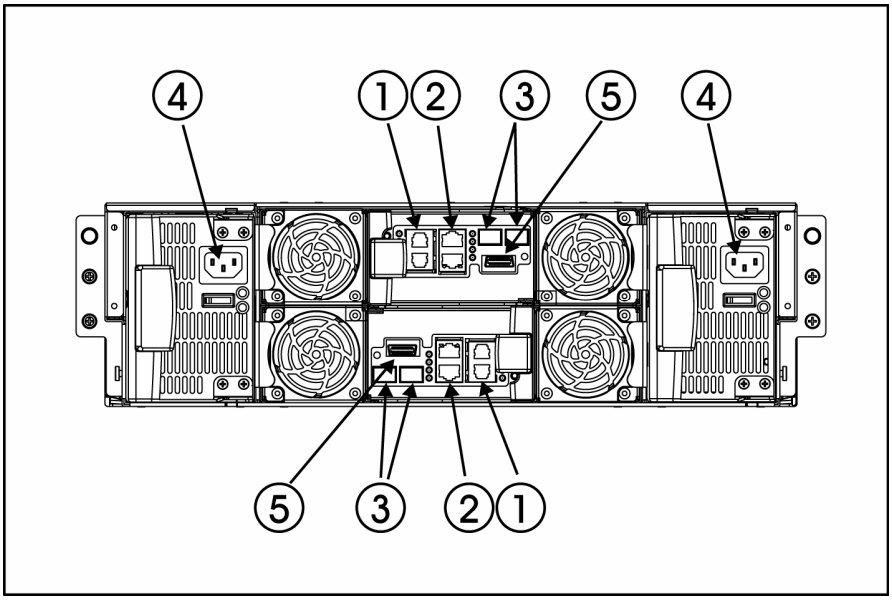

Figure 87 PROVIGO 410 System Connector Locations

Host Interfaces

PROVIGO 410 provides two Gigabit Ethernet interface ports per controller with a total of four available in dual controller configurations (See **Figure 87**).

Teaming Mode

The two Gigabit Ethernet ports on each controller are pre-configured to work in teaming mode. The goal of teaming services is to provide fault tolerance and link aggregation across both Gigabit Ethernet ports. This provides load-balancing, fault-tolerance, and VLAN features across both ports.

Teaming Default

PROVIGO 410 is able to provide all the teaming modes. By default, SLB mode is enabled. Fault tolerance and link aggregation are both supported in this mode. For iSCSI, there is not much advantage to and advised against other teaming modes such as Generic Trunking or 802.3ad.

Power Supply Modules

The enclosure contains one or two power supply modules to provide redundant power for the entire array. One power supply provides enough power to boot up and run a fully loaded system. The second power supply serves as a backup in the event of a system failure.

Each power supply module contains a non-removable fan to cool that power supply module. The module shuts itself down if the specified limit is exceeded for temperature or output current.

Power Supply (LEDs) & On/Standby Switch

Each power supply has a DC On/Standby Switch and two LED status indicators (See **Figure 88**).

- Power Switch – The power supply comes with a two-position switch which controls the DC electrical power to the disk drive modules and other components of the enclosure. The switch must be set to "on" for the power supply module to start. If the Power Switch is on for either power supply module, the enclosure will power up if mainline power is present and at least one power supply module is correctly installed.

- LED #1 – Is a dual LED that includes a green and yellow color indicator. When AC power is applied to the power supply and standby voltages are present the light is yellow. It turns green when all power outputs are available.

- LED #2 – Red if there is a power supply failure

Refer to table below for specific power supply conditions and corresponding LED color combinations.

	Pwr Switch ①	**LED** ②	**LED** ③
No AC Power	Standby	Off	Off
Power Supply Failure	On	Off	Red
Limit Exceeded	On	Off	Red
AC Present	Standby	Yellow	Off
Power Supply DC Output On & Ok.	On	Green	Off

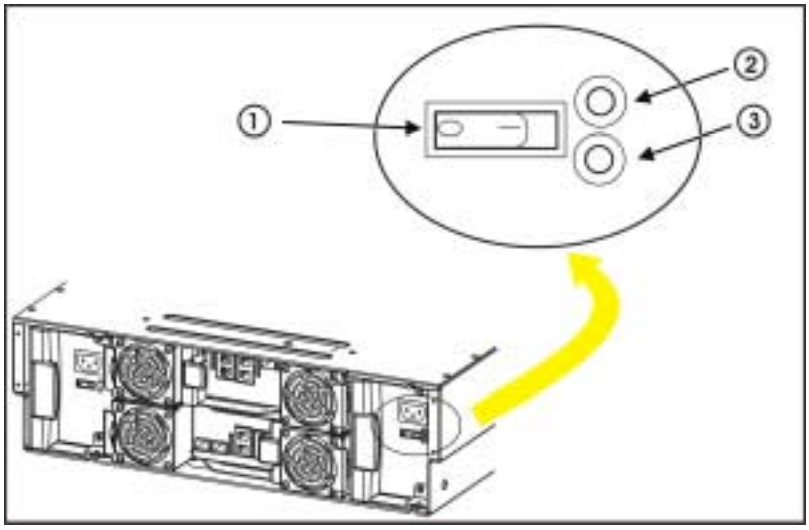

Figure 88 Power Supply Module LEDs and Switch

Exchanging a Redundant Power Supply Module

Note: Unless you have a particular reason, do not switch off the host system or shut down the PROVIGO 410 when removing or installing power supply modules.

1. Go to the back of the enclosure. On the power supply module that you are removing, set the DC Power On/Standby switch ① to "Standby" (See **Figure 88**).

2. Unplug the power cable ② from the power supply module that you are removing (See **Figure 89**).

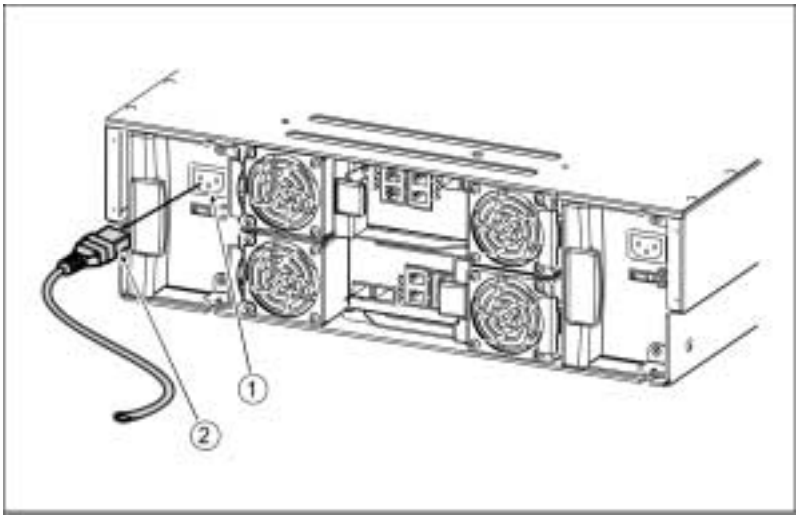

Figure 89 Switching off & Disconnecting the Power Supply Module

Note: The power supply module lights may stay on for a short time. Wait until the AC Power and DC Power lights are off before you go to the next step.

3. Pinch the latch on the power supply module handle ① then pull the handle away from the module as shown (See **Figure 90**).

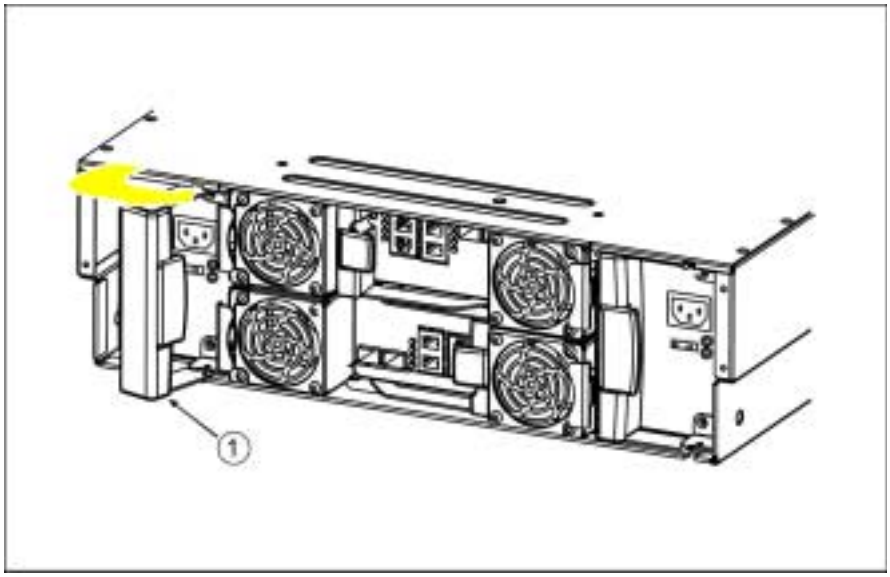

Figure 90 Unlatching the Power Supply Module

4. Pull the power supply module straight out to remove it, supporting the module from below (See **Figure 91**).

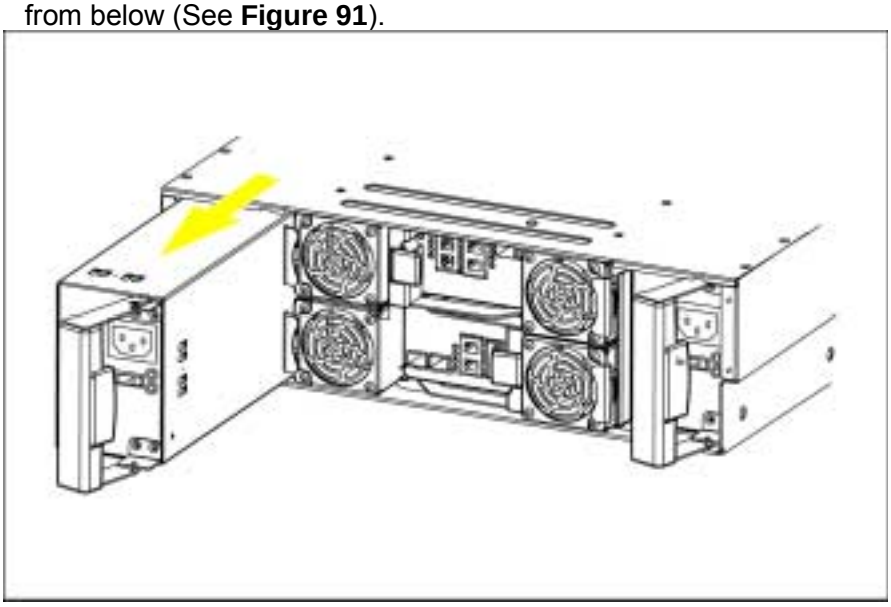

Figure 91 Removing the Power Supply Module

5. Replace parts in the reverse sequence.

Danger: Do not try to open the covers of the power supply module. The power supply module is a replaceable part. Do not try to repair or exchange any component parts.

Caution: Do not insert hands or tools into empty slots.

Note: Remember to set the DC On/Standby switch back to "On" after you have reinstalled the power supply module and power cable. If the DC On/Standby switch is already set to "On" when you reinstall the power supply module, the yellow or red LEDs may come on. If this happens, set the DC On/Standby switch to "Standby", then to "On". The LED should turn green.

Removing All Power from System

It is important to follow the following sequence when removing power from the system to ensure the system does not go into recovery mode. The sequence will ensure the controllers and all components are shut down gracefully.

1. Stop all operations on hosts connected to the controller(s) on the PROVIGO 410.
2. Log-in to the PROVIGO 410's PROVIGO 410 RAIDconn and use the "Shutdown" command at the top right of the screen or use the CLI "dashutdown" command (See "**Appendix B – Using the Command Line Interface**" on page**135**) to stop all activity on the controllers.

NOTE: Make sure that the green Heartbeat LED indicators on both controllers are not blinking. The LED should be either solid green or off before removing power from the system.

3. Set each power supply module Power Switch ① in the enclosure to standby. The green light will turn off.
4. Turn off power from the rack PDU or UPS and remove the power cords ② from each power supply.

5. To return power, turn power on power from the rack PDU or UPS. Reinstall the power cables to each PROVIGO 410 power supply and turn the power switch on for each power supply.

Caution: Before pulling a power cord from a power supply, ensure the LEDS on the power supply are off

Danger:

- Ensure power from rack PDU or UPS is turned off and you shut down the system via PROVIGO 410 RAIDconn before removing power cords from power supplies.
- Never install or reconfigure this product during an electrical storm.
- Do not plug a power cable into the power supply of a PROVIGO 410 unless all modules are completely inserted in their slots and the enclosure is properly installed in the rack with its retaining screws tightened and all host interface cables connected.

UPS Vaulting

UPS Vaulting is provided with the 2.7 release of the KDI. When connected to a smartups device, the storage system will initiate a graceful shutdown when the battery drops below 20% charge or if there are less than 10 minutes remaining runtime available from the battery backup system.

Connection

The COM2 Port on the PROVIGO 410 controllers is used for connecting to a UPS system with a UPS cable 940-1524C from APC or equivalent. The system is approved to work with APC Smart-UPS® systems. No other vendor product has been tested at this time. The UPS cable is provided by the UPS manufacturer.

One cable can be connected to the COM2 port of either controller in a dual controller configuration.

The UPS system must be on and connected to the PROVIGO 410 Storage Array at boot time in order to use the UPS vaulting feature. Upon boot up, the COM2 port is examined

to see if a cable is installed. If a cable is not found, the UPS daemon exits and the feature is disabled.

If a UPS cable is removed while the system is running, the user will be notified via standard email notification. A subsequent email is sent when the cable is restored.

Configuration

No configuration is required or supported. The 20% minimum and 10 minute runtime requirements are preset.

Messages

Messages from the UPS monitoring daemon are stored in /var/log/messages.

During initial boot up, you will find the following entry in /var/log/messages

```
apcusd:  apcupsd startup succeeded.
```

If the UPS is not present, you will find the following entry in /var/log/messages indicating that you are not connected to a UPS. It may be necessary to check the cable or the configuration of the Smart-UPS® device.

```
apcupst[1450]: No UPS found. If one is attached, check the
serial cable and the configurations in DEVICE directive and
UPSCABLE directive.
```

Finally, when there is no UPS present the process shuts down. The following message is found in /var/log/messages.

```
apcusd:  apcupsd error shutdown completed.
```

Notification

The system utilizes email notification system to notify designated individuals of various events in the system. While copies of the notifications are entered into the system logs, you must set up email notification prior to any event to receive these notices through email.

The following is a list of other messages that can be generated and sent to designated recipients via the email notification system.

Event Code	User Message
100	The UPS is stopping power to the Controller.
101	Power from the UPS to the Controller has been stopped.
102	Communication was lost to the UPS. Check the UPS Cable.
103	Communication has been restored to the UPS.
104	Power loss was detected on the UPS.
105	A power failure to the UPS was detected. Controllers now running on batteries.
106	Power has returned to the UPS.
107	Battery power on the UPS is exhausted. Shutting down the Controllers.
108	The battery time limit was exceeded on the UPS. Shutting down the Controllers.
109	The remaining battery charge is below the limit set on the UPS. Shutting down the Controllers.
110	The remaining battery runtime is below the limit set on the UPS. Shutting down the Controllers.
111	The UPS is initiating the reboot sequence.
112	The UPS has initiated the Controller Shutdown Sequence.
113	The UPS is having power problems. Stop your work and log off.
114	Beginning emergency shutdown of Controller because of UPS failure – possibly due to battery failure.
115	Emergency! Batteries have failed on the UPS Change them NOW.

Cache Battery Modules

Each controller contains one battery module to protect the cache in the event of a power failure. The battery module protects the cache for up to 72 hours. A battery module consists of Lithium-ion batteries enclosed in the module.

WARNING - Do not try to replace or exchange the rechargeable battery pack integral to the battery module.

- Keep the battery away from fire.
- Do not expose the battery to water or rain.
- Do not attempt to disassemble the battery.
- Do not short circuit the battery.
- Keep the battery away from children.

Do not put the battery into rubbish that is put into landfill sites. When discarding the battery, observe local ordinances, regulations, and your company's safety standards for disposal.

Hard Disk Modules

PROVIGO 410 can be populated with 15 hard drive modules. Each module consists of SATA or SAS disk drive housed in a hot plug canister with dual ported access to each controller.

Rules for Adding Storage

PROVIGO 410 introduces the concept of disk groups for adding storage to its array. A disk group is composed of an associated group of disks with similar RAID properties. The following are rules that must be followed when creating disk groups for PROVIGO 410:

Disk Group (DG) Parameters	RAID 0 & RAID 5	RAID 1 & RAID 1+0
1. Minimum # of Drives/DG	3	4
2. Drives are Added	In Singles	In Pairs
3. DGs Allowed per System	1-4 *	1-3 *
4. Maximum DG Capacity	6TB **	6TB **
5. RAID Level Allowed/DG	1 ***	1 ***
6. Global Spare Capability	Yes ****	Yes ****
7. Access to DG by both Ctrls	Yes *****	Yes *****

* Drives not assigned to a disk group will be classified as “Global Spares”, Failed Disks or “Orphans” (disks disassociated from their disk group) by the system.

** Capacity of a disk group can be expanded by adding disks from the available pool of global spare disks. Adding disks to a disk group may require multiple disks.

*** Each DG must be assigned a RAID level (0, 1, 5, or 1+0)

**** A Global Spare cannot be added to RAID 0 DG.

***** DGs can be accessed by either controller, but must be assigned a primary controller.

Note: To improve the utilization of the available storage capacity and ensure reliability, it is recommended that all disk drive modules within an enclosure be the same size and from the same manufacturer. This is because for each RAID Group, the system scales down to the size of the smallest drive (except for RAID0).

Installing Disk Drives or Slot Covers

Ensure that the handle ① is fully open on the disk drive module that you are going to install (See Figure 92). With one hand supporting the base of the module and the other hand holding the handle, insert the module and push it into the slot. When the handle touches the front surface of the subsystem, the module will stop. The module is not yet fully inserted. While continuing to push the disk drive module into the slot, slowly close the handle ① until it stops with a click. This action fully inserts the module.

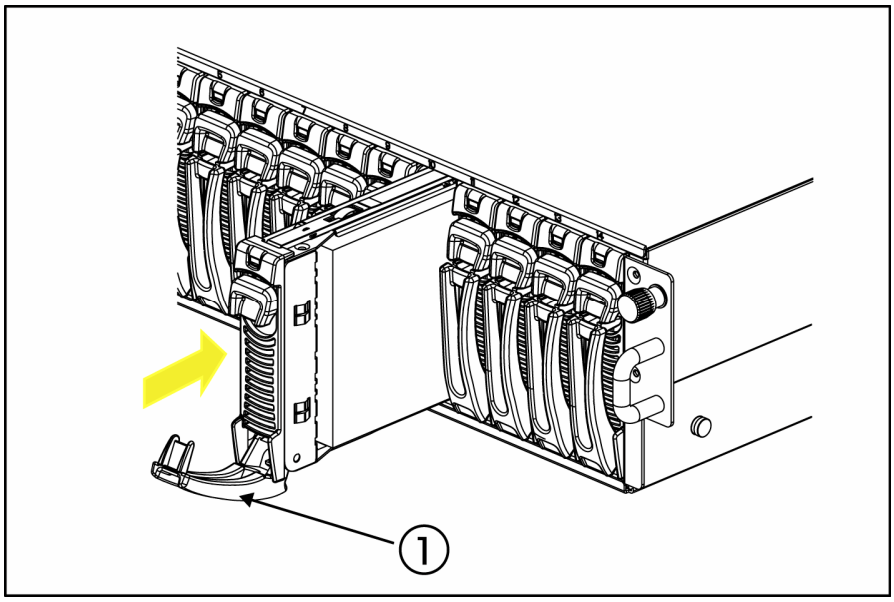

Figure 92 Installing a PROVIGO 410 Hard Disk Module or Slot Cover

Attention: Disk drive modules are fragile. Handle them with care and keep them away from strong magnetic fields. Any slot that has no disk drive module installed must contain a slot cover. The slot cover ensures that the correct airflow is maintained.

Caution: As you push the module into its opening, the lever automatically moves toward the closed position. Ensure that your fingers do not become pinched between the lever and the module.

Verify that the disk drive module that you have just installed is aligned with the sides of the enclosure and that no gap exists between this module and the modules that are next to it. Verify also that the front edge of this disk drive module aligns with the front edges of the modules that are next to it. If the disk drive module is not correctly aligned, remove it and reinstall it. Repeat these steps for each disk drive module you wish to install. Install slot covers that ship with the system in any empty drive slots.

Removing Disk Drives or Slot Covers

Push the latch ① in, and pull the handle ② down. This action pulls the disk drive module partially out of its slot (See **Figure 93**).

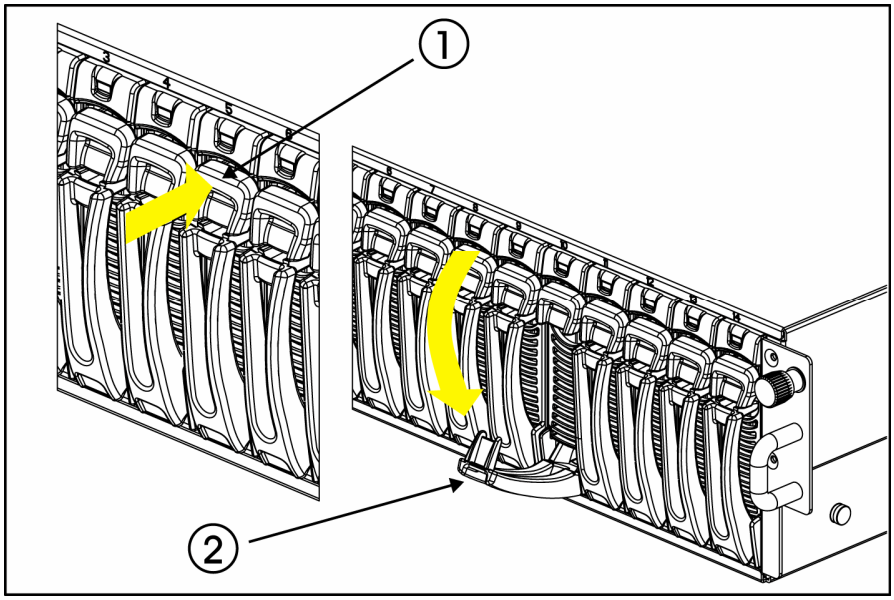

Figure 93 Opening the Handle of a PROVIGO 410 Disk Drive Module

Grip the handle, and carefully pull out the disk drive module (See **Figure 94**). As the module comes out, put one hand under its base to prevent it from falling. Lay the disk drive module on its side. It rests on supports on the bottom of the disk drive module. This prevents the disk drive module from falling over and being damaged.

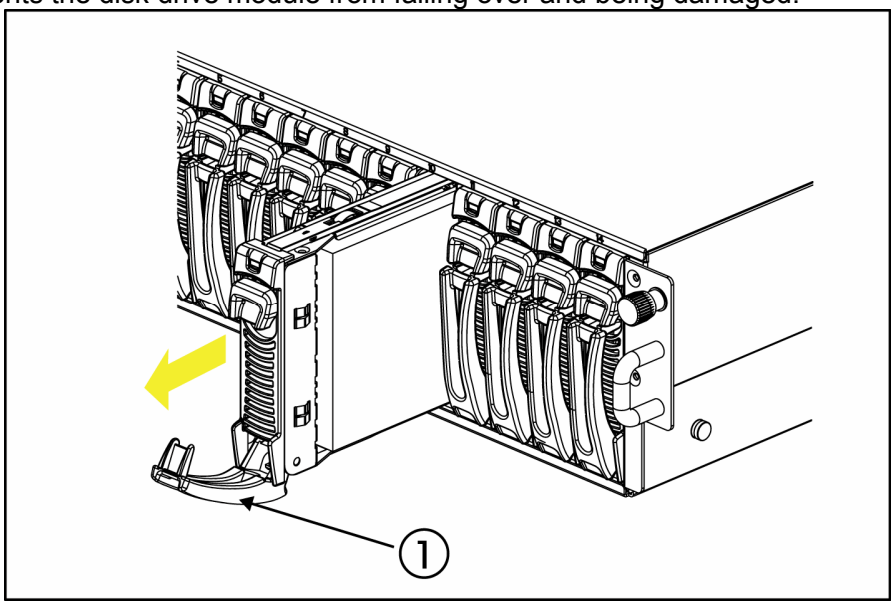

Figure 94 Removing a PROVIGO 410 Disk Drive Module

Caution: Never remove a disk that is part of a disk group during a power failure or improper shutdown. Any information in system cache protected by the battery backup feature will need to be flushed to disk once power is restored. If the system cannot find the appropriate LUN due to an incomplete disk group, the data may be lost.

Hard Drive LEDs

The Hard Disk Modules have a multicolor LED ① at the top of the module (See **Figure 95**). The LED has two color indicators, red and green. Each LED has three states, on, off, and flashing. Refer to table below for specific hard disk module conditions and corresponding LED color combinations.

HDM State	LED
Not Powered	Off
Powered & Idling	Solid Green
Active	Blinking Green
Failed (Replace)	Solid Red
System Health Alert	Blinking Red (All HDMs)

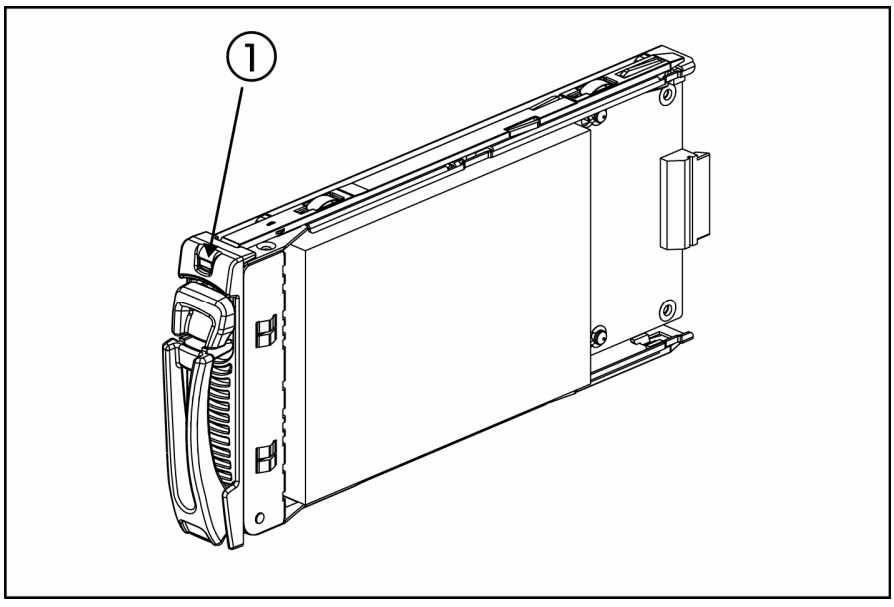

Figure 95 PROVIGO 410 Disk Drive Module Light Location

Disk Roaming

The PROVIGO 410 have the unique capability of allowing disk groups and their properties to roam with the drives. The user can physically remove the drives that make up a disk group and install in different locations within the same array or in a totally different array and the system will recognize the disk group configuration including RAID volume.

You must stop all activity on a Disk Group before attempting to use the Disk Roaming feature. You may either make the Disk Group inactive from the Disk Group properties page or shutdown the system before moving disks from a Disk Group.

Masking and host information are contained in the controller. You will need to reconfigure the masking and host information if you move the disk group to a new controller, as in moving the disk group to a new chassis.

This is possible because when a disk group is created, the system automatically creates a small 128MB partition on each drive. This partition contains the disk group configuration. In this way, all the information regarding a disk group can be recognized by any PROVIGO 410 controller during the automatic boot up discovery or during a manual discovery using the PROVIGO 410 RAIDconn or CLI.

You should delete the Disk Group from the original controller after removing the physical disks. You can not discover the Disk Group when it is reinstalled in the original chassis if you have not removed the Disk Group prior to re-installing it in the original chassis.

The drives can be reinserted in different slots and they can be rediscovered while the system is on-line. You can take advantage of this flexibility to arrange your storage for different on-line and off-line tasks. This feature provides a basic and inexpensive way to upgrade a system and retain your data or move drives to another system for basic system recovery.

Note: All drives in a DG needed for automatic boot up and discovery need to be present for the disk group discovery to take place. A disk group with one missing member can be forced online using the PROVIGO 410 RAIDconn or CLI discovery command.

Disk Orphans

If some of the drives from a disk group are not re-inserted or a drive from the disk group goes bad and is not reconstructed before it roams to a new system, the system will not

be able to discover the disk group. It will know that one or more drives are missing and will designate the drives in the incomplete disk group as orphans. The data in the incomplete set is maintained, however user intervention is required to reset the drives. This ensures that the data on the remaining drives is not overwritten without you intervening. If the drives are designated as orphans, the user has the following options:

Complete the Set – You can track down the remaining drives from the disk group, insert them in any order into the same PROVIGO 410 enclosure and use the PROVIGO 410 RAIDconn to do a discovery. Once the whole set is discovered, the disk group will be restored.

Rebuild a Disk Group – If the original group is part of a RAID 5, 1 or 10 disk group, a bad drive maybe the reason for an orphan classification by the system. If this is the case, make sure you have a global spare in the system and then use the PROVIGO 410 RAIDconn to do forced discovery of the disk group. Once the disk group is discovered and the bad drive is identified, the system will start the RAID rebuild using the global spare. This will restore the disk group.

Re-configure the Drives – If you cannot complete the disk group or rebuild it, then you will not be able to recreate the disk group. transtec recommends to always back up your data to recover from such a scenario. Alternatively you may just want to reuse drives from an old disk group where the data on the drives is no longer needed. In either case you can create a new disk group or turn the drives into global spares by using the PROVIGO 410 RAIDconn.

Appendix A – RAID Concepts & Cache Configurations

RAID Levels Defined

RAID 0 – Data Stripping without Protection - This technique has striping but no redundancy of data. The data is broken down into blocks and each block is written to a separate disk drive.

- Pros - I/O performance is greatly improved by spreading the I/O load across many channels and drives.
- Cons- Not “true" RAID because it is NOT fault-tolerant. The failure of just one drive will result in all data in an array being lost.
- Ideal applications - those that require high bandwidth without need for data protection.

RAID 1 – Disk Mirroring - Consists of at least two drives that duplicate the storage of data. There is no striping.

- Pros- Read performance is improved since either disk can be read at the same time. Write performance is the same as for single disk storage. It provides the best performance and the best fault-tolerance in a multi-user system. 100% redundancy of data means no rebuild is necessary in case of a disk failure, just a copy to the replacement disk. Also, under certain circumstances, RAID 1 can sustain multiple simultaneous drive failures.
- Cons - highest disk overhead of all RAID types (100%).
- Ideal applications - those that require very high availability.

RAID 5 – Distributed Parity - RAID-5 stores parity information but not redundant data. The parity information can be used to reconstruct data. RAID-5 requires at least three disks for the array.

- Pros - This is the best combination of performance and cost.
- Cons – Disk failure has some impact on throughput.
- Ideal Applications -this is the most versatile RAID level and is applicable to a broad range of applications. It is best for multi-user systems in which performance is not critical or which do few write operations.

RAID 10 – Combination of mirroring and striping - This type offers an array of stripes in which each stripe is a RAID-1 array of drives.

- Pros -This offers higher performance than RAID-1
- Cons- but at much higher cost.

- Ideal applications- those requiring high performance and fault tolerance.

Cache Settings Defined

Write Back Cache Option - Data normally written to disk is first written into the cache. During idle machine cycles, the data is written from the cache into disk.

- Pros- Write back cache improves performance because a write to the cache is faster than to disk. Write-back caching yields better performance than write-through caching because it reduces the number of write operations to disk.
- Cons- With this performance improvement comes a risk that data may be lost if the system crashes. The write back cache configuration should only be selected if you have the battery-backed option installed and operating properly on your PROVIGO 410 system.

Write Through Cache Option - Data is written to disk and also written into the cache.

- Pros – This method ensures complete data integrity protection, good performance for sequential access, which relies on less cache. If a subsequent read operation needs that same data, read performance is improved, because the data is already in the cache.
- Cons- Write performance is not improved with this method.

Auto Cache Setting – This setting will automatically turn on the write back cache option and enable the mirrored cache feature when there are two controllers installed. In the event the second controller is removed or fails, the system will automatically switch to the write through cache configuration.

- Pros - Good cache performance and complete data protection.
- Cons – Less optimized performance for sequential access due to mirrored cache system overhead.

Appendix B – Using the Command Line Interface

Preparing to use the CLI

The following information is for experienced user that wants to connect a terminal direct to the Serial port of the PROVIGO 410 or connect via telnet session to the PROVIGO 410 and configure the array using the Command Line Interface (CLI) provided with the system.

This section describes how to access the CLI and provides a list of commands that can be used for system configuration or management and administrative functions of the PROVIGO 410 GUI.

Connecting a Terminal to Access the CLI

Connect the supplied Serial cable to the COM 1 Serial management port on either controller. Connect the opposite end of the cable to a workstation or laptop with a terminal emulation program. A DB9 connector is provided for the workstation or laptop. This is usually connected to the serial comm. 1 port on the workstation or laptop.

Use the following terminal emulation settings to access the CLI:

- Baud Rate – 115,200
- Data Bits – 8
- Parity None
- Stop Bits – 1
- Flow Control Off

Accessing the CLI

Once connected and the PROVIGO 410 is running, and you connect through a terminal emulator to the comm1 port, you will be prompted for the following information:

- Login – root
- Password – root

When connecting through a telnet session, you must log onto the system as **guest**, and use the su command to become root. Most commands must be executed as root.

Using the Network Setup Wizard

The first time you login using root, the system will launch the network setup wizard that will walk you through the network configuration. The default values are in the brackets.

PROVIGO 410 can be configured from either controller. Regardless of which controller is used, the results are the same. After the first-time setup, the PROVIGO 410 system configuration including IP address and hostname are initialized from configuration files stored in the flash of the local controller or from a remote controller. The system configuration of both controllers can be changed at any time via the dasetup and system initialization commands.

When using the network setup wizard, you can press enter to select the default value or type in the private network information you received from your network administrator.

Here is a sample of what you will be asked to provide:

- Enter the IP address for Controller 0 [192.168.0.1]:
- Enter the IP address for Controller 1 [192.168.0.2]:
- Enter the Netmask [255.255.255.0]:
- Enter the Default Gateway [ ]:
- Enter the Name Server [ ]:

When you enter the Name Server IP address, you have finished the network setup.

Note: This wizard will only run automatically when you log in from the Serial terminal as root. It will NOT run when you login using telnet or log in as a non-root user. Once the network is configured, the wizard will not run automatically again. However, you can run the wizard manually by issuing the command 'dasetup'.

.

The default IP address for controller 0 (the lower controller) is 192.168.0.1. The default IP address for controller 1 (the upper controller) is 192.168.0.2. The default netmask is 255.255.255.0.

Using dasetup to change the IP address of the controllers will terminate the current telnet session once you enter the Name Server address. Similarly, if you change the IP address with the network setup window on the GUI, the current session will terminate and you will have to reconnect to the system.

Local/Remote Option

Some of the commands have a local/remote option. This will enable the local controller and remote controller to have different configurations such as the host interface IP addresses. Command *local* means you are executing the command on the local controller, and command *remote* means you are executing the command on the remote controller. The output of these commands will indicate whether the output is from the local controller or the remote controller. The output will also show you the controller id of the local or remote controller.

Some of the commands don't have a local/remote option. Setting the configuration on the local controller with most commands will set the same configuration on the remote controller. The output of these commands will indicate the whether the output is from the local controller or the remote controller. The output will also show you the controller id of the local or remote controller.

Index of CLI Commands

The following table lists the PROVIGO 410 commands available through the command line interface. Any additional executable files that exist in the system that may be called by these commands, but should not be executed as standalone commands.

command	description
daacl	used to manage access control list
daaddrbk	used to add email address to notification list
dacomment	used to add comments to configuration
daconnstat	used to display connection statistics
dadate	used to display or set time and date
dadg	used to manage disk groups
dafcnode	used to display or set fibre channel node name
dagiveback	used to return IP address after failover or upgrade
dahelp	used to display information about commands
dahost	used to define and manage host
dahostname	Allows user to set hostname for storage array
dahwstat	used to display hardware status
daifconfig	Used to display or configure network interface
dainspread	used to manage Inspection Read utility
dainventory	Used to inventory the hardware components
daisns	used to manage isns registration
dakem	used to display event log
dalun	used to manage lun masking table
dampio	used to enable or disable MPIO

dans	used to display or set name server
dapasswd	used to change user password
dareboot	used to reboot local or remote controller
daroute	used to display or set gateway address
dasendlog	used to send system logs
dasetlogmail	used to define addresses for sendlog
dasetup	used to set IP addresses, netmask, gateway and name server
dashutdown	used to shutdown local or remote controller
dasmart	used to run SMART commands
dasnmp	used to manage snmp groups
dastat	used to display system information
dastonith	used to force reboot of remote controller
dasysalert	used to enable or disable blinking lights
datarget	used to set target name
daum	used to upgrade or mange kernel images
davd	used to manage virtual disks
daversion	used to display versions of kernel images or boot images

Commands

Command line interface commands can be grouped into four general categories. The following sections provide a list of commands used for system administration, network administration, storage administration or support activities.

System Administration Commands

daacl

Use daacl to manage the access control list on at a system level.

```
daacl      - manage access control list.

Usage:
    daacl [Options]

Options:
     -h/--help:                       help
     -L/--list:                       list all allowed addresses
     -a [address/prefixlen]:          add an address into the allowed list
     -d [address/prefixlen]:          delete an address from the allowed list

Command Examples:
     daacl -a 0.0.0.0/0:              allow access from all
     daacl -a 192.168.1.0/24:         allow access from one subnet
     daacl -a 192.168.1.1/32:         allow access from one IP
```

dacomment

Use dacomment to add or display comments on a system.

```
dacomment     - add comments to configuration.

Usage:
    dacomment [Options]
    -h:                         help
    -g:                         get the comment
    -s ["comments"]:            set the comment

Examples:
    dacomment -s "this is storage system for \"server A\""
    dacomment -g
```

dadate

Use dadate to display or set the system date and time on one or two controllers. Without any option specified, the current date is printed out.

```
dadate     - Use dadate to display or set the system date and time on
             one or two controllers

Usage:
    dadate [Options]

Options:
    -help:                          help information
    -s [MMDDhhmmYY.ss]:             set date and time specified
    -z [GMT|gmt|UTC|utc] offset:  set time zone by offset to GMT/UTC

Command Examples:
    dadate:                       prints out the current date and time
    dadate -s 1216183005.32:      sets the date and time
    dadate -z GMT:                set as Greenwich Mean Time
    dadate -z GMT +8:             set as 8th zone east to GMT
    dadate -z GMT -8:             set as 8th zone west to GMT

Notes:
For offset value for your location, please check out the following sites
    http://www.greenwichmeantime.com
    http://www.worldtimezone.com
```

dahelp

Use dahelp to display a list of CLI commands

```
dahelp - display command usage page.
```

```
Usage:
    dahelp [Options]

Options:
    -h/help:         display this help page
    [command]:       display usage for the specified command
    -s [pattern]:    display usage for all the commands that match the pattern

Command List:
daacl      - manage access control list.
daaddrbk   - add email address to notification list.
dacomment  - add comments to configuration.
daconnstat - display connection statistics.
dadate     - display or set the date and time.
dadg       - manage disk groups.
dagiveback - return IP address after failover or upgrade.
dahelp     - display help information about commands.
dahost     - define or manage host.
dahostname - display or set hostname for the storage array.
dahwstat   - display hardware status.
daifconfig - display or configure network interface.
dakem      - display event log.
dalun      - manage the lun masking table.
dampio     - enable or disable MPIO.
dans       - display or set the name server.
dapasswd   - set the password for the user.
dareboot   - reboot local or remote controller.
daroute    - display or set gateway address.
dasendlog  - send system logs.
dasetlogmail - define the email addr for dasendlog.
dasetup    - set IP addresses, netmask, gateway and name server.
dashutdown - shutdown local or remote controller.
dasnmp     - Add/Delete SNMP managers to send traps to.
dastat     - display the system information.
dastonith  - force reboot of remote controller.
dasysalert - enable or disable blinking lights.
datarget   - set target name.
daum       - upgrade or manage kernel images.
davd       - manage virtual disks.
daversion  - display versions of kernel images and boot image.
```

dahost

Use dahost to add or delete host entries from the system.

```
dahost     - define or manage host.

Usage:
    dahost [Options] [Parameter Options]

Options:
    -a:                     add a new entry
    -d:                     del an entry
    -l:                     print out all the entries
```

```
Parameter Options:
    -b [type]               0 is iSCSI, 1 is FC
    -w [host's WWN]:        world wide name
    -u [user name]:         empty by default
    -o [host OS type]:      0 is Windows by default, 1 is Linux
    -n [host name]:
    -i [host ip address]:   empty by default (default is Dynamic)
    -p [initiator secret]:  empty by default
    -s [target secret]:     empty by default
    -t ["comments"]:        empty by default. Use quotes for strings

Option Requirements:
    option a requires:      w and n are mandatory. i,u,p,s and t is optional
    option d requires:      w or n

Command Examples:
    dahost -a -w iqn.test3 -o 0 -n test3:     add a host
    dahost -d -w iqn.test3:                   delete a host
    dahost -d -n test3:                       delete a host
```

dahostname

Use the dahostname to change the host name for the system. For those systems that build an iSCSI target name with the host name, you must use the –f (force) flag to change the hostname because target names may be affected by changing the host name.

```
dahostname     - display or set hostname for the storage array.

Usage:
    dahostname [-h | -f] [hostname]

Options:
    -h/help:            display usage
    -f:                 force hostname change. This flag must be specified to
                        change host name for system supporting multiple
                        iSCSI targets.

Command Examples:
    dahostname:              Display current hostname
    dahostname -f newhost:   Change hostname to newhost
```

dapasswd

Use dapasswd to change the passwords on the system.

```
dapasswd     - set the password for the user

Usage:
    dapasswd [Options] [Option Parameters]
Options
```

```
    -help:                  display help page
    -p:                     new password
Option Parameters:
    Username:               specify username. Default is current user

Command Examples:
    dapasswd -p passwd username: sets password for a user
    dapasswd -p passwd:          sets new password for current user
```

dareboot

Use dareboot to reboot the local or remote controller. If no controllers are specified, dareboot will reboot the local controller.

```
dareboot     - reboot local or remote controller.

Usage:
    dareboot
    dareboot [Options]

Options:
    -help:              help
    -local:             reboot local controller
    -remote:            reboot remote controller

Command Examples:
    dareboot -help
    dareboot
    dareboot -local
    dareboot -remote

Notes:
The "dareboot" without any option has the same effect as
"dareboot -local".
```

dashutdown

Use dashutdown to shutdown the local or remote controller. If local or remote are not specified, dashutdown operates on the local controller.

```
dashutdown   - shutdown local or remote controller.

usage:
    dashutdown [Options] [Option parameters]

Options:
    local:              shutdown local controller.
                        Default is local controller
    remote:             shutdown remote controller

Option Parameters:
    -t [seconds]:       execute at now plus seconds specified
    now:                shutdown now
```

```
Notes:
       dashutdown requires a time be specified, either “now” or –t <seconds>.
```

dasnmp

Use dasnmp list, add or delete SNMP managers

```
dasnmp     - Add/Delete SNMP managers to send traps to.

Usage:
    dasnmp [Options]

Options:
    -h/--help:                 help
    -L/--list:                 list all SNMP managers
    -a/--add [address]:        add a SNMP manager with/without community
    -d/--delete [address]:     delete a SNMP manager
    -c [community string]:     community string of a SNMP manager

Option Requirements:
    option c requires:         a

Command Examples:
    dasnmp -a 192.168.1.1:     add a SNMP manager without community
    dasnmp -a 10.0.0.1 -c pub: add a SNMP manager with community "pub"
    dasnmp -d 10.0.0.1:        delete a SNMP manager
```

dastonith

Use dastonith to force a reboot on the remote controller by applying a low pulse to the reset circuit.

```
dastonith  - force reboot of remote controller

Usage:
    dastonith [-h/help] [0/1]
Options:
    -h/help:          print this help page.
    1:                stonith remote controller. This is same as press
                      and hold on the remote controller's reset button.
    0:                revive the remote controller.  This is same as
                      release the remote controller's reset button.
Command Examples:
    dastonith 1
    dastonith 0

Notes:
To stonith the remote controller, issue the dastonith
commands in the following sequence.

    dastonith 1
```

```
    dastonith 0

The system will reboot after you issue dastonith 0.
```

dasysalert

Use dasysalert to display system alert setting or to change system alert setting to on or off. There are two components in the system alert mechanism. Alert LEDs on the hard drives blink 3 times over 2 second intervals when there is a system alert, and alert messages in the system health message system. Disabling system alerts only turns off the blinking LEDs. This command does not affect the health messages viewable on the system health page or in system logs.

```
dasysalert     - turn on/off system alert

Usage:
    dasysalert [-h] [on | off]


Options:
    -h:                   help and print this page
    -s/-show:             the current status, 'on' or 'off'
    On:                   turn on system alert
    Off:                  turn off system alert

Command Examples:
    dasysalert on
    dasysalert off
    dasysalert -show
```

Network Administration Commands

dafcnode

Use dafcnode to display or set the fibre channel node name

```
dafcnode - display/set fibre channel node name

Usage:
    dafcnode [-h] [-g] [-s <fibre_channel_node_name>]

Options:
    -h/help:                        help text
    -g:                             get the fibre channel node name
    -s [fibre_channel_node_name]:   set fibre channel node name

Command Examples
    dafcnode -g
    dafcnode -s a111234567890120
```

daifconfig

Use daifconfig to manage the network interface of the controllers on the system.

```
daifconfig - display or configure network interface.

Usage:
    daifconfig [Options] [Option parameters]

Options:
    local | remote:          specifies host interface location.
                             Default is local.

Option Parameters:
    <interface>:             should be da0
    <IPAddress>:             IP address
    <netMask>:               netmask for this subnet

Command Examples:
    daifconfig local da0 192.168.0.1 netmask 255.255.255.0
    daifconfig remote da0 192.168.2.1 netmask 255.255.255.0
```

daisns

Use daisns to manage isns registration

```
daisns     - register with isns server.

Usage:

    daisns [Options] [parameters]

Options:
    -h/--help:        help page
    -l/--list:        list iSNS server IP and alias name
    --start:          start iSNS service
    --stop:           stop iSNS service
    --setserver:      save iSNS server ip and alias name
    --display:        display the iSNS server online or offline

Parameters:
    --addr:           iSNS server IP address
    --alias:          iSNS server alias name

Option Requirements:
    option setserver: addr and alias are mandatory.

Comand Examples:
   daisns --start
   daisns --stop
   daisns -l
   daisns --stop
   daisns --display
   daisns --setserver --alias test3 --addr 10.4.1.3
```

dans

Use dans to list, add or delete a name server IP address to the network configuration.

```
dans     - display or set the name server.

Usage:
    dans [-h | -l | -a IPAddress | -d IPAddress]

Options:
    -h:                   help
    -l:                   list all name servers
    -a:                   add a name server
    -d:                   delete a name server

Command Examples:
    dans -a 192.168.123.5; Add a new name server
    dans -d 192.168.123.5; Delete a name server
```

daroute

Use daroute to list the route table, to add a default gateway or remove a default gateway from the controller network configuration.

```
daroute     - display or set gateway address.

Usage:
    daroute [Options] [{add | del} default gw <gwIPAddr>]

Options:
    -h:                   help

Option Parameters:
    add:                  add a default gateway
    del:                  delete default gateway
    <gwIPAddr>:           gateway IP address to add or delete
Command Examples:
    daroute add default gw 192.168.123.254
```

dasetup

Use dasetup to perform initial system configuration from the console. dasetup is a wizard that will prompt the user for network configuration information. The following shows the usage of the dasetup wizard:

```
This wizard will walk you through the network configuration. You can configure
the rest of the system using the web GUI. The default values are in the
bracket, you can press "Enter" to select the default value.  If you make any
mistake, you can press "ctrl-c", then logout and login to run this wizard
again.

Enter IP address for Controller 0 [192.168.0.1]:
Enter IP address for Controller 1 [192.168.0.2]:
```

```
Enter netmask [255.255.255.0]:
Enter IP address for Default Gateway [ ]:
Enter IP address for Name Server [ ]:

You have finished configuring the network setup.
If you want to reconfigure the network, you can run dasetup again.
Now you can configure the rest of the system using the web GUI at
http://192.168.0.1.
```

Storage Administration Commands

dadg

Use dadg to create and manage disk groups.

```
dadg     - manage disk groups.

Usage:
    dadg [Options] [Parameter Options] [Device List]

Options:
    -h:                         help
    -help:                      help
    -Ldisk:                     list disk information
    -Ldg [disk-group name]:     list disk-group information
    -Lall:                      list generic information for all disk groups
    -a [disk-group name]:       create a disk group
    -d [disk-group name]:       delete a disk group
    -D [disk-group name]:       deactivate a disk group
    -A [disk-group name]:       activate a disk group
    -r:                         discover new disk groups
    -e [disk-group name]:       expand a disk group
    -c [disk-group name]:       change a disk-group parameter
    -y [single device name]:    recycle an orphan or faulty disk to a spare
    -EC [disk-group name]:      check and correct RAID consistency
Parameter Options:
    -o [controller owner]:      [0, 1] for controller 0, controller 1
    -l [RAID level]:            [0, 1, 5, 10] for RAID-0,1,5,10
    -v [alarm level]:           [1-99] for disk-space alarm threshold
    -n [new disk-group name]:   new disk-group name
    -m [comment string]:        explanatory comment
    -f:                         force deletion of disk group with virtual disks,
                                force activation of degraded disk group,
                                force discovery to find degraded disk groups,
                                force name change of disk group
Option Requirements:
    option a requires:          o, l, v, and device list. m is optional
    option d requires:          f is optional
    option A requires:          o. f is optional
    option r requires:          o. f is optional
    option e requires:          device list
    option c requires:          at least one of v, n, or m. f is required with n
    option y requires:          one device only
Device List Mapping:
    /dev/xda:                   corresponds to slot 0 (the leftmost slot)
    /dev/xdb - /dev/xdn:        slot 1 - slot 13
    /dev/xdo:                   slot 14 (the rightmost slot)
```

```
Device List Requirements:
    RAID-0, RAID-5 require:   3 to 15 devices in device list
    RAID-1, RAID-10 require:  4, 6, 8, 10, 12, or 14 devices in device list
Device List Examples:
    /dev/xda /dev/xdb /dev/xdc /dev/xdg /dev/xdh /dev/xdo
    /dev/xd[a-c,g-h,o]
    /dev/xd[a,b,c,g-h,o]
Command Examples:
    dadg -help
    dadg -Ldisk
    dadg -Ldg dg1
    dadg -Lall
    dadg -a dg1 -o 0 -l 5 -v 95 -m "comment" /dev/xd[a-e]
    dadg -d dg1
    dadg -d dg1 -f
    dadg -D dg1
    dadg -A dg1 -o 1
    dadg -A dg1 -o 1 -f
    dadg -r -o 0
    dadg -r -o 0 -f
    dadg -e dg1 /dev/xd[f-i]
    dadg -c dg1 -m "comment new" -v 90
    dadg -c dg1 -n dg1new -f
    dadg -y /dev/xdo
    dadg -EC dg1
```

dagiveback

Use dagiveback when the system is in Active/Takeover mode to return the system to Active/Active state. The system must be in Active/Active state to provide Takeover capability in the event of a controller failure.

dagiveback must be run on the Active controller.

```
dagiveback     - return IP address after failover or upgrade.

Usage:
    dagiveback
    dagiveback -help
```

dalun

Use dalun to present or un-present virtual disk to host.

```
dalun      - manage the lun masking table

Usage:
    dalun [options] [fields...]

Options:
    -h:              help
    -a:              add a new lunmask entry
    -d:              del a lunmask entry
    -p:              print out all the existing lunmask entries

Option Requirements:
    -n {host]
```

```
    -g [disk_group_name]
    -k [virtual_disk_name]
    -P [parent_virtual_disk_name]

Notes:
The -P parent_virtual_disk_name is only required if the
virtual disk presented is a snapshot virtual disk.


Option Notes:
      -u [LUN_ID]:      Specifies LUN ID

Command Examples:
    dalun -a -n host1 -g dg1 -k vd1:
                  presents host1 to virtual disk vd1 in the
                  disk group dg1

    dalun -d -n host1 -g dg1 -k vd1:
                  un-present host1 from virtual disk vd1 in the
                  disk group dg1.

    dalun -d -n host2 -g dg1 -k snap1 -P vd1:
                  un-present host2 from snapshot snap1 of virtual
                  disk vd1 in the disk group dg1
```

dampio

Use dampio to enable or disable MPIO on the system.

```
dampio     - enable or disable MPIO.

Usage:
    dampio [Options]

Options:
    -help:              help
    -enable:            enable MPIO
    -disable:           disable MPIO
    -show:              show current MPIO setup

Command Examples:
    dampio -help
    dampio -enable
    dampio -disable
    dampio -show
```

datarget

This is not used in releases that have different iSCSI target names for each lun. Use display or set iSCSI target names for a controller.

```
datarget     - set target name.
```

```
Usage:
    datarget [Options]

Options:
    -help:              help
    -set [IQN name]:    setup iSCSI target IQN names
    -show:              show iSCSI target IQN names
    -reset:             reset iSCSI target IQN names to factory default

Command Examples:
    datarget -help
    datarget -set iqn.2005-05.com.mycompany:myarray1
    datarget -show
    datarget -reset
```

davd

Use davd to create and manage virtual disk.

```
davd     - manage virtual disks.

Usage:
davd [Options] [Parameter Options]

Options:
    -h:                         help
    -help:                      help
    -L [virtual-disk name]:     list virtual-disk information
    -a [virtual-disk name]:     create a data virtual disk
    -S [virtual-disk name]:     create a snapshot virtual disk
    -d [virtual-disk name]:     delete a virtual disk
    -e [virtual-disk name]:     expand a virtual disk (data only)
    -c [virtual-disk name]:     change a virtual-disk parameter
Parameter Options:
    -g [disk-group name]:       disk group containing the virtual disk
    -s [size]:                  size in MB.
                                Append G(or g) or T(or t) for GB or TB
    -m [cache policy]:          [0, 1, 2] for auto, write-through, write-back
    -r [read ahead policy]:     [0, 1] for enable, disable
    -n [virtual-disk name]:     new virtual-disk name
    -o [0/1]:                   I/O type of a data vd:
                                block i/o (0, default), file i/o (1)
    -p [parent virtual disk]:   data virtual disk associated with a snapshot
    -i [initial space]:         [1-100] for initial space allocated to a snapshot
    -w [warning level]:         [0-100] for snapshot warning level
    -f:                         force deletion of virtual disk or snapshot
                                force name change of virtual disk or snapshot
    -P [0/1]:                   share type: 0 for public, 1 for private
    -W [0/1]:                   access type of share:
                                0(default) for writable, 1 for read-only
    -u [share users]:           a list of users delimited by ',' for private share

                                prefix a user name with '+' or nothing
                                  for adding a user to private share,
                                prefix a user name with '-'
                                  for deleting a user from private share

Option Requirements:
    option L requires:          g, (p if snapshot)
    option a requires:          g, s, m (P if -o 1)
```

```
    option S requires:        g, p, P (i and w are optional)
    option d requires:        g, f, (p if snapshot)
    option e requires:        g, s
    option c requires:        g, (p if snapshot), (n,f) or m or [P, y, u]
Option Notes:
    option e:                 Expansion is only allowed for data virtual disks,
                              -s [size] specifies the capacity increment
    option c:                 Name changes of VD and snapshot (the -n option)
                              require the -f option,
                              -m [cache policy] is only allowed for data VDs

Command Examples:
    davd -help
    davd -L lv1 -g dg1
    davd -L snap1 -g dg1 -p lv1
    davd -a lv1 -g dg1 -s 10240 -m 2 -r 0
    davd -a lv1 -g dg1 -s 10G -m 2
    davd -a lv2 -g dg1 -s 1g -m 1 -o 1 -P 0
    davd -S snap1 -g dg1 -p lv1
    davd -S snap2 -g dg1 -p lv2 -i 20 -w 85 -P 1 -u user1,+user2
    davd -d lv1 -g dg1 -f
    davd -d snap1 -g dg1 -p lv1 -f
    davd -e lv1 -g dg1 -s 8192
    davd -e lv1 -g dg1 -s 8G
    davd -c lv1 -g dg1 -m 1
    davd -c lv2 -g dg1 -W 1
    davd -c lv2 -g dg1 -P 1
    davd -c lv2 -g dg1 -u +user1,-user2,+user3,user4
    davd -c lv1 -g dg1 -r 1
    davd -c lv1 -g dg1 -n lv1new -f
    davd -c snap1 -g dg1 -p lv1 -n snap1new -f
    davd -c snap2 -g dg1 -p lv2 -u user3,-user1
    davd -c snap2 -g dg1 -p lv2 -P 0
    davd -c snap2 -g dg1 -p lv2 -n snap2new -f
```

System Support and Miscellaneous Commands

daaddrbk

Use daddrbk to add email addresses to the notification list. Email addresses included in the notification list receive email on system traps.

```
daaddrbk     - add email address to notification list.

Usage:
    daaddrbk [-hL] [-ad <email_address>]
Options:
    -a [email_address] :        add an email address to the address book
    -d [email_address] :        delete an email address from the address book
    -h/help:                    help
    -L:                         list all the current email address

Command Examples:
    daaddrbk -a steve@foo.com: add a notification email address
    daaddrbk -d steve@foo.com: delete a notification email address.

Note:
```

```
To receive email notification, you need to make sure the
default gateway/router and name server is setup properly.
The system needs to be able to talk to the mail server of
the email address you enter here.
```

daconnstat

Use daconnstat to display iSCSI or Fiber Channel statistics.

```
daconnstat     - display connection statistics.

Usage:

    daconnstat [Options]

Options:
    -help:              help
    -show:              show statistics
    -reset sid:         reset rx/tx counters of session "sid" to zero

Command Examples:
    daconnstat -help
    daconnstat
    daconnstat -show
    daconnstat -reset 0x7bc13000

Notes:
The "daconnstat" without any option has the same effect as
"daconnstat -show".
```

The following is a sample of the output from daconnstat:

```
iSCSI statistics:
<Controller 0>
--
SID:       0x7bc0c000
Initiator: iqn.1991-05.com.microsoft:rack10
Target:    iqn.2000-08.com.intransa:30-31.dg1.vd1
Rx Data:   0 B
Tx Data:   18 MB
--

<Controller 1>
--
SID:       0x7bc1e000
Initiator: iqn.1991-05.com.microsoft:rack10
Target:    iqn.2000-08.com.intransa:30-31.dg1.vd2
Rx Data:   0 B
Tx Data:   1634 B
--
```

dahwstat

Use dahwstat to show information of chassis components, drives and alert messages.

```
dahwstat     - show hardware status

Usage:
    dahwstat [Options]

Options:
    -help:              display help information
    -drives | -drive:   show properties of all the drives
    -drive [no.]:       show properties of a specified drive
    -chassis:           show information of chassis components, including
                        controller ID, second controller installed,
                        Second controller presence, fan readings,
                        temperature readings, voltage readings and battery
                        information
    -alert              show alert messages if any

Command Examples:
    dahwstat -alert
    dahwstat -drive 0
    dahwstat -drives
    dahwstat -chassis
```

Examples Output from commands:

```
   # dahwstat -alert
              [12/21 12:11:28]      Disk Group(s) degraded or unhealthy

  # dahwstat -drive 0
                Drive:                            Drive0 (/dev/sda)
                Model:                                   ST380013AS
                Serial Number:                           3JV3J6BX
                Firmware Revision:                       3.05
                Size [KB] (1 KB = 1024 Bytes):           78150744 KB
                Size [GB] (1 GB = 1024^3 Bytes):         74.5 GB
                Size [GB] (1 GB = 10^9 Bytes):           80.0 GB
                Read Total  [MB] (1 MB = 1024^2 Bytes):  0.0 MB
                Write Total [MB] (1 MB = 1024^2 Bytes):  0.0 MB
                Read block requests:                     6
                Write block requests:                    0

   # dahwstat -drives    -- You will see properties of a list of all the drives

   # dahwstat -chassis
                Controller ID:                 0
                Second controller installed:   1
                Second controller presence:    1
                Power Supply information:
                   [PS0]: Present, AC Good, DC Good
                   [PS1]: Present, AC Good, DC Good
                Fan speed readings: (in RPM)
                   [Fan tray 1 - Fan 0]: 5004 (Normal)
```

```
      [Fan tray 1 - Fan 1]: 5702 (Normal)
      [Fan tray 0 - Fan 0]: 5042 (Normal)
      [Fan tray 0 - Fan 1]: 5678 (Normal)
   Temperature readings: (in Celsius degree)
      [T1 - CPU temperature]:          41 (Normal)
      [T2 - North bridge temperature]: 47 (Normal)
      [T3 - System temperature]:       49 (Normal)
   Voltage readings: (in volt)
      [V1]: 2.5 (Normal)
      [V2]: 1.5 (Normal)
      [V3]: 3.3 (Normal)
      [V4]: 5.0 (Normal)
      [V5]: 12.2 (Normal)
   Battery Information:
      Remaining capacity percentage: 100 %
      Estimated backup time:         81.9 Hours
```

dainspread

Use dainspread to manage and run the Inspection Read utility.

```
dainspread - manage Inspection Read.

Usage:
  dainspread [Options] [Option Parameters]

Options:
  -h/-help:  Show help information
  -start:    Immediately start Inspection Read
  -stop:     Immediately stop Inspection Read
  -resume:   Immediately start Inspection Read from last inspected block
  -schedule: Add, delete or show Inspection Read schedule. When no parameters
             are provided, -schedule shows the current schedule
  -progress: Show the progress of Inspection Read

Option Parameters:
  -repeat:           Run continuously until dainspread is issued with the stop
                     option or the specified -duration is satisfied.
  -cpu [priority]:   Set the Inspection Read priority based upon cpu load [1-3].
                     High is 1, Medium is 2, Low is 3. The priority is low
                     unless otherwise specified.
  -duration [hours]: The maximum of hours to run [1 - 10000]. Inspection Read
                     will stop after running for the specified number of hours.
  -add:              Add a new schedule or change the current schedule.
  -del:              Delete the current schedule.
  -sch_resume:       Schedule Inspection Read to resume from last inspected
                     block.
  -day [Day]:        Used to specify the day of week to start Inspection Read
                     [0 - 6], where 0 specifies Sunday.
  -hour [hour]:      Specifies the hour to start Inspection Read [0 - 23], where
                     0 specifies midnight.

Option Requirements:
  start:             repeat, cpu and duration are optional.
  stop:              no parameters should be specified with stop.
  resume:            repeat, cpu and duration are optional.
  schedule:          add or del required. If add, then day and hour are
                     required. sch_resume, repeat, cpu and duration are
                     optional. If del, then no other parameters should be
                     specified.
```

```
  progress:          no parameters should be specified with progress.

Examples:
  dainspread -start
      Start Inspection Read.
  dainspread -start -repeat
      Start Inspection Read and run continuously until dainspread is issued with
      the -stop option.
  dainspread -start -repeat -cpu 2 -duration 6
      Start Inspection Read at a medium priority and repeat until 6 hours have
      passed.
  dainspread -stop
      Stop Inspection Read
  dainspread -resume
      Resume Inspection Read from the last inspected block.
  dainspread -resume -cpu 1
      Resume Inspection Read from the last inspected block at a high priority.
  dainspread -resume -repeat -cpu 2 -duration 6
      Resume Inspection Read at a medium priority and repeat until Inspection
      Read has run for 6 hours.

  dainspread -schedule
      Show the current Inspection Read schedule.
  dainspread -schedule -add -day 0 -day 1 -hour 23
      Schedule Inspection Read to start every Sunday and Monday at 11 PM.
  dainspread -schedule -add -day 0 -hour 23 -duration 6
      Schedule Inspection Read to start every Sunday at 11 PM and run for 6
      hours.
  dainspread -schedule -add -day 0 -hour 23 -duration 6 -sch_resume
      Schedule Inspection Read to resume from the last inspected block every
      Sunday at 11 PM and run for 6 hours.
  dainspread -schedule -add -day 0 -hour 23 -repeat -cpu 2 -duration 6
      Schedule Inspection Read to start every Sunday at 11 PM at a medium
      priority and repeat until Inspection Read has run for 6 hours.
  dainspread -schedule -del
      Delete the current Inspection Read schedule.

  dainspread -progress
      Show the progress of the current Inspection Read.
```

dainventory

Use dainventory to create a detailed inventory of the system in a text file with the chassis serial number, date and time the inventory is taken. Use ftp to extract the file /pub/chassisinfo from the system. This file does not remain after reboots because the controller may be moved between different chassis.

```
      dainventory   - Creates /pub/chassisinfo text file
                      that contains a description of the
                      system components.

                      This file is available for ftp transfer
                      with an anonymous ftp connection.

                      You may enter a chassis serial number as a parameter
                      a parameter or enter the seial number when running
                      the command
```

```
Usage:
   dainventory [Options] [Parameters]

Options:
   -h              help
   -help           help

Parameters:
    Chassis Serial number

Examples:
   dainventory
   dainventory TY0304A1006
   dainventory -h
```

dainventory -help

dakem

Use dakem to display the event log.

```
dakem      - display event log.

Usage:
    dakem [Options] [Option parameters]

Options:[-hvEi] [-s severity] [-C n]

    -h :          help.
    -v :          version.
    -E :          display KEM messages.
    -i :          suppress the special char display.

Option Parameters:
    -s [severity]:    [info][warning][critical][all(default)]
    -C [n]:           number of messages to be displayed, maximum of 6000

Command Examples
    dakem -E:                       Display the last 256 KEM messages.
    dakem -E -C 6000:               Display all the KEM messages.
    dakem -E -C 6000 -s warning:    Display all the KEM message of severity
                                    warning and above.
```

dasendlog

Use dasendlog to send the system log files from both controllers to the recipient as defined with dasetlogmail.

```
dasendlog     - send system logs.

Usage:
    dasendlog [Options]

Options:
    -s:           send the log file to the tech support
    -h/-help:     help page
```

dasetlogmail

Use dasetlogmail to set the email address that is used when you click on the send logs hyperlink in the GUI help page. This email address is not the address used with the notification functionality.

```
dasetlogmail     - define the email addr for dasendlog.

Usage:
    dasetlogmail [Option] email_address
Option:
    -h/--help:                  help

Command Example:
    dasetlogmail abcd@yahoo.com: system log will be sent to abcd@yahoo.com
```

dasmart

Use dasmart to run the SMART (Self-Monitoring, Analysis & Reporting Technology) commands on specific disks in the system.

```
dasmart - run SMART commands

Usage:
    dasmart [Options] [Device]

Options:
    -help:                  print help page
    -i | -identify:         identify drive and SMART supported/enabled status
    -e | -enable:           issue SMART ENABLE OPERATIONS command
    -d | -disable:          issue SMART DISABLE OPERATIONS command
    -n | -enable_autosave   issue SMART ENABLE ATTRIBUTE AUTOSAVE command
    -v | -disable_autosave  issue SMART DISABLE ATTRIBUTE AUTOSAVE command
    -c | -capabilities:     issue SMART READ DATA to read capabilities
    -h | -healthstatus:     issue SMART RETURN STATUS to read health of drive
    -a | -attributes:       issue SMART READ DATA to read attributes of drive
    -r | -errorlog:         issue SMART READ LOG to read the error log
    -m | -immedtest:        run SMART off-line immediate test
    -b | -immedtest_abort:  abort SMART off-line immediate test
    -f | -immedtest_status: read status of SMART off-line immediate test
    -s | -selftest_short:   run SMART short self-test in off-line mode
    -x | -selftest_ext:     run SMART extended self-test in off-line mode
    -g | -selftest_abort:   abort the SMART self-test currently running
    -u | -selftest_status:  read status of SMART self-test currently running
    -t | -selftest_log:     issue SMART READ LOG to read the self-test log
    -o | -online_status:    report if software driver considers drive on-line
Device Mapping:
    /dev/xda:               corresponds to drive in slot 0 (the leftmost slot)
    /dev/xdb - /dev/xdn:    slot 1 - slot 13
```

dastat

Use dastat to display system status.

```
dastat     - display the system information.
```

```
Usage:
    dastat
    dastat -help
```

The following is an example of the output from dastat.

```
------------------------------------------------------------
Name:                                 30-31-0
Controller Slot#:                     0
Remote Controller:                    Up
Operational State:                    Active/Active
Total Capacity:                       934.25 GB
Time:                                 Thu Nov 17 15:13:07 UTC 2005
Kernel Version (local):               K:2.1.47-20050811
Kernel Version (remote):              K:2.1.47-20050811
CPLD Board Midplane Rev (local):      C:5 B:1.0 M:2
CPLD Board Midplane Rev (remote):     C:5 B:A04 M:2
BIOS Version (local):                 04300S01 12/14/04
BIOS Version (remote):                04301S01 12/21/04
Fibre Channel WWN:                    200000a0d1e16350

   Drive    View from controller 0    View from controller 1
     sda .............IN DG ("dg1") ................NOT OWNER
     sdb .............IN DG ("dg1") ................NOT OWNER
     sdc .............IN DG ("dg1") ................NOT OWNER
     sdd .............IN DG ("dg1") ................NOT OWNER
     sde ..................MISSING ..................MISSING
     sdf ..................MISSING ..................MISSING
     sdg ..................MISSING ..................MISSING
     sdh ..................MISSING ..................MISSING
     sdi ..................MISSING ..................MISSING
     sdj ..................MISSING ..................MISSING
     sdk ..................MISSING ..................MISSING
     sdl ..................MISSING ..................MISSING
     sdm ..................MISSING ..................MISSING
     sdn ..................MISSING ..................MISSING
     sdo ..................MISSING ..................MISSING


No System Health Alert
------------------------------------------------------------
```

daum

Use daum (upgrading utility) to upgrade the system image

There are two system images on each controller. An image pointer is used to indicate which one to use for next boot operation. With daum, the user can manually synchronize the two images and choose which one to use for next boot operation.

The new image file must be located on a FTP site and daum uses a URL to obtain the new image file.

This command must be run as super user.

```
daum       - upgrade or manage kernel images

Usage:
    daum [Options]

Options:
    -h/--help:              help information
    -i/--info:              display information about the system images,
    -s/--switch:            switch the image pointer for next boot operation
    -y/--sync:              synchronize current and backup images
    -n [URL]/--new [URL]:   upgrade with a new system image
    -v                      show version information of daum

Option Requirements:

    URL is in the following format:
    usrname:passwd@ip_address/path_name/file_name

    The host specified with 'ip_address' must be a FTP site. If the FTP
    server accepts 'anonymous' user, the 'usrname:passwd' can be neglected.

Command Examples
    daum -n joe:passwd@192.168.1.1/images/newkdi
    daum -n 192.168.1.1/pub/kdi.0.1
```

daversion

Use daversion to display the version of the boot image and kdi versions. Without any option specified, the version information of local controller is printed out.

```
daversion - displays the version of boot image and kdi versions.

Usage:
    daversion [Options]

Options
    -h/--help:       help information
    Local:           print out version information of local controller
    Remote:          print out version information of remote controller

Command Examples:
    daversion -help
    daversion
    daversion local
    daversion remote
```

The following is an example of the output from daversion.

```
daversion
```

```
********** local diskarray controller ID 0 **********************
Current RedBoot image release Version: 3.00-20051003
Current RedBoot image release Date: Mon Oct  3 19:37:48 2005

Current KDI image release Version: 2.1.208-20051109
Current KDI image release Date: Wed Nov  9 21:20:24 2005

Backup KDI image release Version: 2.1.200-20051026
Backup KDI image release Date: Wed Oct 26 17:54:34 2005

*****************************************************************
```

Appendix C – System Specifications

This section describes the physical characteristics of the PROVIGO 410 and its environmental and power requirements.

System Enclosure Dimensions

Height	Width	Depth
131.2mm (5.2in.)	445mm without ears (17.5in.)	629mm (24.76in.)

Weight

Fully Loaded System	Without Drives
90 Lbs.	60 Lbs.

The weight of a PROVIGO 410 depends on its configuration.

Rack Service Clearances

Minimum Rack Service Clearance		
Front	Back	Sides
114cm (45in)	81cm (32in)	No service clearance required on sides

Environment

Temperature and Humidity

	Air Temperature	Relative Humidity	Maximum Wet Bulb
Operating	10°C to 30°C (50°F to 86°F)	10% to 90% non-condensing	27°C (80°F)
Recommended operating range	20°C to 25°C (68°F to 77°F)	20% to 50%	
Non-operating	10°C to 50°C (50°F to 122°F)	10% to 80% non-condensing	27°C (80°F)
Storing	1°C to 60°C (20°F to 60°F)	5% to 80% non-condensing	29°C (84°F)
Shipping	−40°C to 60°C (−40°F to 140°F)	5% to 100% condensing but not precipitating	29°C (84°F)

Altitude (from sea level)

Operating:

- 0 to 3,000 m (with the maximum environment temperature degrade in 1.8°C/1,000m)
- 0 to 10,000 ft (with the maximum environment temperature degrade in 1.0°F/1,000ft)

Non-operating/Storage/Shipping:

- 0 to 10,000 m
- 0 to 30,000 ft

Note- *The upper dry bulb temperature limit of the rack-mounted unit is lowered by 1.8 °C (1.0 °F) for every 1000 meters (1000 feet) above 0 meters (0 feet). The upper wet bulb temperature limit is lowered by 1.8 °C (1.0 °F) for every 1000 meters (1,000 feet) above 0 meters (0 feet).*

Airflow

PROVIGO 410 system requires approximately 2.8 m^3/min. (100 ft^3/min.) of airflow.

Electrical Power

Electrical power is supplied to each PROVIGO 410 unit by one or two power supply modules. These power supply modules convert the input voltage to dc for distribution within the enclosure. Each disk drive module receives power from one or both power supply modules; however, one power supply module provides enough power for a fully populated enclosure. If one power supply module fails, all hard drive modules can be initiated and continue to operate.

Input-Voltage Requirements

The power supply modules need a minimum of 100-120 VAC at 7A or 200-240 VAC at 3.5A. They can run uninterrupted with ac inputs from 90 to 264 volts and from 47 to 63 Hz. Power line disturbances (within the limits given in the following table) do not interrupt the operation of a PROVIGO 410 unit.

Transient Voltage	Duration	Frequency	Nominal Voltage
264V rms	2.0 s	47 to 63 Hz	240 V rms
90 V rms	2.0 s	47 to 63 Hz	100 V rms

Power Input

The following table shows the power input that is required for a fully populated unit in a PROVIGO 410. The start-up power assumes that no data transfer is occurring as the disk drive modules start up.

Description	Startup	Operating	Two power supplies each
100 V ac input	530 W / 550 VA	520 W / 540 VA	270 W / 280 VA
240 V ac input	520 W / 550 VA	510 W / 540 VA	260 W / 280 VA

Power Factor

Power-factor correction is applied within the power supply modules.

Output Protection

Each power supply module has over-current and over-voltage protection.

Early Power-Off Warning (EPOW)

Each power supply module provides an Early Power-Off Warning (EPOW) signal to the controller module in the system enclosure. When both power supply modules signal an EPOW, the control module will enter into battery backup mode and then gracefully shut down the hard disk modules.

Power Cables

Each unit in a PROVIGO 410 is supplied with a country-specific power cable for each power supply module that it contains.

Power Control

There is a DC On/Standby switch on each power supply module.

Power Sequencing

The disk drive motors start under the control of the control module when the enclosure powers on. All the disk drive modules start their power sequence when the input voltages are inside the specified limits. The other modules in an enclosure have no power sequencing.

Auto Restart

A PROVIGO 410 will automatically restart when input power is restored within specification after a power failure. The disk drive modules' motors start under the control of the controller module.

Acoustic Noise

The A-weighted sound power level measured for a rack mounted unit does not exceed 68 dba. The level was measured under the following conditions.

- 15 disk drive modules installed in a PROVIGO 410.
- The PROVIGO 410 is in rack M/T 9306-M200 (EIA position 15).
- The remainder of the rack was empty.

Safety Approvals

The PROVIGO 410 is certified to the following safety standards:

- Listed by Underwriters Laboratories Inc. in accordance with UL 60950 and CAN/CSA C22.2 60950.
- System IB-3100

- Control board IB-3100-C

Electromagnetic Compatibility - Emissions

- Australia / New Zealand AS/NZS 3548 (Class A)
- European Union EN55022: 1994 (Class A) & 89/336/EEC
- Japan VCCI (Class A)
- Taiwan BSMI CNS13438
- USA Title 47 CFR, Part 15 (Class A)

Electromagnetic Compatibility - Immunity

- European Union EN55024: 1998

Appendix D – Hardware & Software Compatibility List

This section lists the current 3rd party hardware and software tested and supported by transtec.

Supported Operating Systems

Client Server Systems connecting to the PROVIGO 410 can be running any of the following operating systems:

- Windows 2000 & 2003 Server
- Windows XP Professional
- Red Hat Linux 7.3, 8.0, 9.0, ES

Supported iSCSI Initiators

- MS iSCSI Initiator (Up to Version 1.06)
- Open Source iSCSI Linux - formerly Cisco iSCSI Initiator - (Up to Version 3.6.2)
- Intel Pro/1000MT Dual Port Server Adapter
- Intel Pro/1000MT Desktop
- Intel Pro/1000XT Server Adapter
- Intel Pro/1000 T IP Storage Adapter (iSCSI)
- Broadcom NetXtreme Gigabit Ethernet

Supported FC HBAs

LSI HBAs

- LSI7102XP-LC
- LIS7202XP-LC

QLogic HBAs

- QLA2310-F-E
- QLA2340-E-SP
- QLA2342-E-SP

Appendix E PROVIGO 410 Illustrated Parts List

	PN	photo
FIBER_CONTROLLER	PV41F-CTRL-B	
NON-FIBER_CONTROLLER	PV41E-CTRL-B	
Middle Plane Module	PV41-MIDPL-B	

Power Supply	PV41-PSU-B	
Empty SATA2 HDD Module	PV41-TRAY-B	
BATTERY, 0801	PV41-BBU-B	

FAN MODULE	PV41-FAN-B	
DAUGHTER CARD	PV41-DCRD-B	

A

Access RAIDConn Storage Manager, 38
Accessing the CLI, 134
Acoustic Noise, 163
Active/- Mode, 110
Active/Active Mode, 108
Active/Takeover Mode, 108
Active-Active Controller, 16
Airflow, 162
Altitude, 161
Assembling Rail Kit, 32
Audible Alerts, 21
Auto Restart, 163

B

Battery Fault, 103
Battery Modules, 124

C

Cache Auto Setting, 133
Cache Configurations, 132
Cache Policy, 66
AUTO, 29
write-back (WB), 29
write-through (WT), 29
Cache Settings Defined, 133
Change password, 83
Changing GUI Refresh rate, 83
Changing System Health Alert Option, 84
Changing the sytetem password, 83
Chassis, 95
CLI, 39
Accesing CLI, 134
Command list. *See*
daacl, 137
daaddrbk, 151
dacomment, 138
daconnstat, 151
dadate, 138
dadg, 146
dafcnode, 143
dagiveback, 147
dahost, 139
dahostname, 140
dahwstat, 152
daifconfig, 144
dainspread, 153
dainventory, 155
daisns, 144
dakem, 156
dalun, 148
dampio, 148
dans, 145
dapasswd, 140
dareboot, 140
daroute, 145
dasendlog, 156
dasetlogmail, 157
dasetup, 145
dashutdown, 141
dasmart, 157
dasnmp, 142
dastat, 158
dastonith, 142
dasysalert, 143
datarget, 149
daum, 158
davd, 149
daversion, 159
Network Setup, 135
Storage Administration Commands, 146
Support Commands, 151
System Administration Commands, 137
CLI Commands, 136
Command Line Interface, 134
Compatible Hardware & Software, 165
Configuring the System, 38
Connecting a Terminal to Access the CLI, 134
Connecting the Power Cords, 36
Connectors, 115
Console
Connection, 134
Controller Features, 106
Controller Giveback, 111
Controller LEDs, 107
Controller Module Features, 106
Controller States, 108
Converting Orphan to Spare, 61
CTL FAIL, 107

D

daacl, 137
daaddrbk, 151
dacomment, 138
daconnstat, 151
dadate, 138
dadg, 146
dafcnode, 143
dagiveback, 147
dahelp, 138
dahost, 139

dahostname, 140
dahwstat, 152
daifconfig, 144
dainspread, 153
dainventory, 155
daisns, 144
dakem, 156
dalun, 148
dampio, 148
dans, 145
dapasswd, 140
dareboot, 141
daroute, 145
dasendlog, 156
dasetlogmail, 157
dasetup, 145
dashutdown, 141
dasmart, 157
dasnmp, 142
dastat, 158
dastonith, 142
dasysalert, 143
datarget, 149
daum, 158
davd, 149
daversion, 159
DC On/Standby Switch, 117
Dedicated SATA Channels, 17
Degraded System Availability, 102
disk drive LEDs, 100
Disk Group
Controller Selection, 54
Creation, *52*
Creation Rules, 52
Discovery, 60
Disks per RAID level, 53
Expanding, 61
Manual Give Back, 114
Minimum number of drives, *52*
Number in a chassis, 28
Orphan, *60*
Orphan to Spare, 60
Orphaned, 59
Properties Page, 55
RAID 0, 52
RAID 1, 52
RAID 10, 52
RAID 5, 52
RAID Build Progress, 59
RAID Rules, 53
Roamiing, 130
Roaming, 59
Selecting RAID Level, *53*
Disk Group Properties, 50
Disk Group Setup, 50
Disk Groups and Capacity, 28
Disk Orphans, 130
Disk Roaming, 17, 130
DNS server information, 30
Drive Status Indicators (LEDs), 129

E

Early Power-Off Warning, 163
ECC Memory Error, 103
Electrical Power, 162
Electromagnetic Compatibility - Emissions, 164
Electromagnetic Compatibility - Immunity, 164
Electrostatic Discharge, 12
Email Notification, 92
Environment Specifications, 161
Ethernet Activity, 107
Ethernet connector, 115
Ethernet Teaming, 23, 26
Event
Battery Charge Fault, 103
Battery Fault, 103
Downgrade, 102
Fan Failure, 102
Memory Error, 103
Network Connection Failure, 104
Temerature Exceeded, 101
Unhealthy RAID, 103
Voltage Exceeded, 102
Exchanging a Redundant Power Supply Module, 118

F

Fail-back of Disk Groups, 111
Fan Failure, 102
Faulty Battery Charger, 103
Fibre Channel connector, 115
Field Replaceable Modules, 19

G

Gateway address, 30
Global Spares, 17
GUI
Banner Area, 43
Login Screen, 42
Navigation Tabs, 44
organization, 42
Setup Wizard, 44
Status Icon, 46
Work Area, 45

H

HB Link, 107
HDM State, 129
Health Alerts, 100
Host
Properties, 78
setup, 75
Host Interfaces, 116
Hot Swap Modules, 105

I

Illustrated Parts List, 167
Indicator Lights, 107
Input-Voltage Requirements, 162
Inspection Read, 89
Progress, 90, 91, 92
rules, 89
Starting, 90
Installation
Disk Module, 35
Installing Chassis in Rack, 34
Installing Disk Drives or Slot Covers, 126
Installing Hardware Modules, 35
Installing the Hardware, 30
IP address, 38
IP Address, 30, 39
configure, 38
default, 38
iSCSI Network Considerations, 25
iSNS configuration, 86

J

Jumbo Frames, 25

K

Kit Contents, 30

L

LED, 129
Local/Remote Option, 136
login information, 40
LUN Masking, 18

M

Maintenance, 95
Manual "Giveback" of IP Address, 110
Mirrored Cache, 16
MPIO, 22
MPIO Option, 85
Multi-Path, 17

N

Net mask address, 30
Network Connection Lost, 104
Network Setup, 47

O

On-Line Firmware Update, 19
Options
GUI Tab, 82

P

Password, 83
Password management, 83
Planning your Installation, 22
Point in Time Snapshot, 18
Power Cables, 163
Power Connector, 115, 117
Power Control, 163
Power Factor, 163
Power Input, 162
Power Output Protection, 163
Power Sequencing, 163
power supply conditions, 117
Pre-configured IP address, 38
Presenting VD, 79
PSU AC PWR light, 117
PSU DC PWR light, 117
PSU Power Switch, 117
PSU Status Indicators (LEDs), 117

R

Rack Preparation, 31
Rack Service Clearances, 161
Rack Warnings, 12
RAID
Build Progress, 59
RAID 0 – Data Stripping without Protection, 132
RAID 1 – Disk Mirroring, 132
RAID 10 – Combination of mirroring and striping, 132
RAID 5 – Distributed Parity, 132
RAID Concepts, 132
RAID guidelines, 125
RAID Levels Defined, 132
RAID Rules, 53
Rail Kit Preparation and Installation, 31
Read Ahead Policy, 67

Rebuild a Disk Group, 131
Re-configure the Drives, 131
Removing All Power From System, 121
Removing Disk Drives, 127
Replacing a Failed Controller (Redundant System, 112
Replacing a Failed Controller (Non-redundant System, 114
Rules for Adding Storage, 125

S

Safety Approvals, 163
Serial port connector (RS232 port), 115
Setting Date and Time, 49
Setup Wizard, 44, 46
Snapshot, 70
Expanding, 74
Expansion, 71, 74
naming, 74
Overflow, 71
Properties, 72
Size, 71
Software
Upgrade, 87
Software RAID, 17
Start Up
Initial Delay, 40
Status
Event Page, 98
Hardware Status Display, 97
iSCSI Statistics, 98
System Logs, 98
Status Indicators, 21
Storage Related Commands, 146
Supported, 165
Supported FC HBAs, 166
Supported Operating Systems, 165
Supported SAS Hard Drives, 165
Symbols on Equipment, 13
System, 104
System Description, 15
System Enclosure Dimensions, 161
System Health Alert, 100
System Initialization, 48
System Log Notifications, 21
System Management, 18
System Options, 82
System Properties, 95
System Specifications, 161
System Upgrade, 19
System Upgrade from GUI, 87
System Warnings, 13

T

Teaming Default, 116
Teaming Mode, 116
Technician Notes, 12
Temperature Threshold Exceeded, 101
TRAY0, 107
TRAY1, 107
Triggers and Corrective Action, 101
Turning Off the System Health Alert, 104

U

Unhealthy RAID Disk Group, 103
Upgrade
Software, 87
Upgrading from Single controller to Dual controller, 114
UPS, 122
calbe, 122
Connection, 122
Messages, 123
Notification, 123

V

Virtual Disk
Cache Policy, 66
Create, 65
Extending capacity, 70
Presenting to Host, 79
Properties, 64
read ahead policy, 67
Setup, 63
Snapshot, 70
Virtual Disks
Maximum Number, 63
Voltage Threshold Exceeded, 102

W

Weight, 161
Write Back Cache Option, 133
Write Through Cache Option, 133